



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

**Worldwide Consortium for the Grid (W2COG)
Research Initiative Phase 1
Final Report**

by

Christopher Gunderson

31 March 2006

Approved for public release; distribution is unlimited

Prepared for:

Office of the Secretary of Defense, Deputy Undersecretary for Advanced Systems and Concepts, Assistant Secretary of Defense of Network Information Integration, Office of Force Transformation, and the Defense Advanced Research Project Office

THIS PAGE INTENTIONALLY LEFT BLANK

NAVAL POSTGRADUATE SCHOOL
Monterey, California 93943-5000

COL. David A. Smarsh, USAF
Acting President

Leonard A. Ferrari
Provost

This report was prepared for and funded by Office of the Secretary of Defense,
Deputy Undersecretary for Advanced Systems and Concepts, Assistant Secretary
of Defense of Network Information Integration, Office of Force Transformation,
and the Defense Advanced Research Project Office

Reproduction of all or part of this report is authorized.

This report was prepared by:

Christopher Gunderson
Co-Principal Investigator

Peter Denning
Co-Principal Investigator

Reviewed and Released by:

Dan C. Boger
Interim Associate Provost and
Dean of Research

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 31 March 2006	3. REPORT TYPE AND DATES COVERED Technical Report	
4. TITLE AND SUBTITLE: Worldwide Consortium for the Grid (W2COG) Research Initiative Phase 1 Final Report			5. FUNDING NUMBERS DWAM50072	
6. AUTHOR(S) Christopher Gunderson				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Cebrowski Institute Naval Postgraduate School, 833 Dyer Road, Sp-537, Monterey, CA 93943			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of the Secretary of Defense, Deputy Undersecretary for Advanced Systems and Concepts, Assistant Secretary of Defense of Network Information Integration, Office of Force Transformation, and the Defense Advanced Research Project Office			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this report are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) OSD provided \$1.6M, and allowed eighteen months, for the W2COG research initiative to learn and apply lessons from the world of "open" e-business on the worldwide web to accelerate GIG development. Compared to a typical DoD Think Tank "study", W2COG more than returned value of OSD's investment by delivering a number of successful process pilots for: 1. Rapid (30-60 day), low cost (10s of \$K), objective, expert, industry analysis of net-ready issues; 2. Community of interest (COI) for "semantic data strategy"; 3. Rapid demonstration, validation, and fielding of bundled interoperable "net-ready" edge-of-the-GIG network components. These pilots performed by the members of a functioning community of international government and industry experts proved the hypothesis that an "open" e-business approach can team mutually motivated government and industry partners in ventures to find accelerated "good enough" paths to GIG functionality. Recommendations are for W2COG sponsors and their constituents to harvest the benefit of their success by using W2COG institute; honest broker, to evaluate, validate and verify capability; risk mitigation and low-cost, rapid turnaround on concepts, pilots, and prototypes; certify and deploy reference implementations via consumable model; "netcentric productivity metrics"; semantic data strategy.				
14. SUBJECT TERMS NCO, netcentric, network centric, GIG, SOA, W2COG, transformation, Cebrowski Institute, FORCEnet, LANDWARNET, C2CONSTELLATION, C2, C4, C4ISR, NCES, NCO/W RM, Data Strategy, IA, e-business, open consortium			15. NUMBER OF PAGES 317	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

OSD provided \$1.6M, and allowed eighteen months, for the W2COG research initiative to learn and apply lessons from the world of “open” e-business on the worldwide web to accelerate GIG development. Compared to a typical DoD Think Tank “study”, W2COG more than returned value of OSD’s investment by delivering a number of successful process pilots for: 1. Rapid (30-60day), low cost (10s of \$K), objective, expert, industry analysis of net-ready issues; 2. Community of interest (COI) for “semantic data strategy”; 3. Rapid demonstration, validation, and fielding of bundled interoperable “net-ready” edge-of-the-GIG network components. These pilots performed by the members of a functioning community of international government and industry experts proved the hypothesis that an “open” e-business approach can team mutually motivated government and industry partners in ventures to find accelerated “good enough” paths to GIG functionality. Obvious recommendations are for W2COG sponsors and their constituents to harvest the benefit of their successful by establishing a governmental process that exploits the W2COG as follows:

- a. Single point of contact for immediate access to a spectrum of broad cross-industry, government, and academia information processing expertise
- b. Honest broker (e.g. Underwriters’ Lab), per status as 501(c)3 tax-exempt charitable/scientific activity, to evaluate, validate, and verify information processing capability
- c. Risk mitigation and low-cost, rapid turnaround on information processing concepts, pilots, and prototypes
- d. Means to certify and immediately field successful information processing reference implementations broadly via consumable off-the-shelf model
- e. Means to establish “netcentric productivity metrics” per operational COI information processing performance targets.
- f. Means to incrementally develop pragmatic semantic data strategy.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

BACKGROUND AND OPERATION CONSTRUCT	1
W2COG OUTCOMES.....	2
RECOMMENDATIONS TO OSD	5
APPENDICES	6
1. ADF IPV6 TRANSITION STRATEGY	7
INTRODUCTION	9
TRANSITION STRATEGY ANALYSIS	18
RECOMMENDED IPV6 TRANSITION STRATEGY	41
IPV6 ADDRESS SPACE REQUIREMENTS	51
IPV6 TRANSITION GOVERNANCE	59
ADO IPV6 WORKFORCE REQUIREMENTS.....	67
RISK MANAGEMENT.....	74
DEPENDENCIES AND KEY ASSUMPTIONS.....	76
CONCLUSIONS	77
RECOMMENDATIONS	80
ANNEX A INTEROPERABILITY OPTIONS.....	81
TRANSLATION	88
ANNEX B PHASE 1 DETAILED PLANNING	90
ANNEX C RISK LOG	92
ANNEX D DCP PROJECT SUMMARY	100
ANNEX E IPV4	101
ANNEX F CIOG ORGANSATIONAL CHART.....	102
ANNEX G MOBILE IP	103
MOBILITY IN IPV6	103
GENERAL MIPV6 BENEFITS	104
MIPV6 STATUS	104
GENERAL MIPV6 ISSUES.....	104
NETWORK MOBILITY.....	105
2. GES SOA STANDARD REVIEW	106
MEMORANDUM FOR DIRECTOR DEFENSE INFORMATION SYSTEM AGENCY	107
EXECUTIVE SUMMARY	109
SUCCESS CRITERIA.....	109
ENDORSEMENT OF XML BASE PROTOCOL STANDARDS.....	109
BENEFITS AND CAVEATS OF A STANDARDS-BASED APPROACH	110
PROPOSED REQUIREMENTS TAXONOMY	111
AP ANALYSIS	112
REFERENCE ARCHITECTURE.....	114
TESTING AND VALIDATION.....	115
CALL TO ACTION	115
OMG COMMENT ON PROPOSED STANDARDS FOR IMPLEMENTING GIG ENTERPRISE SERVICES.....	119
STATEMENT OF WORK FOR REVIEW OF DISA DRAFT STANDARDS FOR IMPLEMENTING GIG ENTERPRISE SERVICES.....	120
3. NORTHCOM WIRELESS STUDY TASK ORDER	121

4. TOWARDS A RICH SEMANTIC MODEL OF TRACK: ESSENTIAL FOUNDATION FOR INFORMATION SHARING.....	124
SUMMARY	125
BACKGROUND	126
SEMANTICS AND PRAGMATICS OF TRACK, BY EXAMPLE.....	130
AN INITIAL SPECIFICATION OF TRACK SEMANTICS	135
USING THE TRACK MODEL TO ACHIEVE THE STATED OBJECTIVES	141
THE R&D AGENDA TO ACHIEVE THE POTENTIAL BENEFITS	143
RELATED RESEARCH AND TECHNOLOGY.....	145
CONCLUSIONS	146
5. MODEL-BASED COMMUNICATION NETWORKS AND VIRT: FILTERING INFORMATION BY VALUE TO IMPROVE COLLABORATIVE DECISION-MAKING	147
ABSTRACT	148
INTRODUCTION AND OVERVIEW	148
THE BASIC PROBLEM WITH CURRENT APPROACHES: STATELESS NETWORKING.....	149
VIRT IMPROVES TIME-STRESSED COLLABORATIVE DECISION-MAKING	151
OVERALL TECHNICAL STRATEGY AND HIGH-LEVEL ARCHITECTURE.....	153
PRODUCT-LINE, COMPONENT-BASED TECHNICAL ARCHITECTURE.....	157
RELATED RESEARCH	166
PRINCIPAL REMAINING CHALLENGES.....	167
NEAR-TERM EXPLOITATION OPPORTUNITIES	168
SUMMARY AND CONCLUSION.....	169
REFERENCES.....	169
6. TWO THEORIES OF PROCESS DESIGN FOR INFORMATION SUPERIORITY: SMART PULL VS. SMART PUSH.....	171
ABSTRACT	172
BACKGROUND	173
PROPOSED APPROACHES	174
ANALYSIS	176
MONITORING CONDITIONS OF INTEREST	183
DISCUSSION.....	186
CONCLUSIONS	188
REFERENCES.....	189
7. VALUABLE INFORMATION AT THE RIGHT TIME (VIRT) TEAM MISSION STATEMENT	191
THE PROBLEM	192
THE OPPORTUNITY	193
OBJECTIVES & GOALS	194
BENEFITS.....	195
TECHNICAL WORK REQUIRED	196
MODELS, SEMANTICS, INFERENCE.....	196
8. REFERENCE IMPLEMENTATION DOCUMENTATION OF HUMANITARIAN PRIVACY-PRESERVING COLLABORATIVE NETWORK	198
HP NetTop™ DEMONSTRATION	199
HP NetTop™ DEMONSTRATION CONCEPT.....	199
HP NetTop™ REFERENCE INSTALLATION	200
TWO ENCLAVE HP NetTop™ INSTALLATION	201
9. BUSINESS TRANSFORMATION AGENCY ENTERPRISE SOFTWARE INCUBATOR PROPOSAL	203

W2COG BACKGROUND AND OPERATION CONSTRUCT:	205
W2COG INSTITUTE VALUE PROPOSITION TO BUSINESS TRANSFORMATION AGENCY	206
10. REFERENCE IMPLEMENTATION DOCUMENTATION OF GIG-LITE ON-LINE ENVIRONMENT	208
TAB A: EQUIPMENT, DATA, AND INTERFACES.....	211
TAB B: HOW TO PUBLISH TO A GIGCAST CHANNEL.....	217
TAB C: MISSION SCENARIOS AND THREADS.....	221
YOU ARE THERE WORKSHOP: HUMANITARIAN DISASTER RELIEF	221
YOU ARE THERE WORKSHOP: BORDER CONTROL	221
11. DOCUMENTATION OF TRUSTED AUTHORIZATION (ROLE BASED) POLICY ENGINE	222
12. NETCENTRIC CERTIFICATION OFFICE STATEMENT OF WORK.....	233
13. SELECTED MEDIA CLIPPINGS	237
CONSORTIUM JUMP STARTS NETWORK-CENTRIC INTEROPERABILITY	238
NETCENTRIC WARFARE LOOKS LIKE COMMERCIAL E-BUSINESS. JUST ASK AL QAEDA.	240
DoD TURNS TO INDUSTRY FOR THE INTERNET IT WANTS	242
14. VALUABLE INFORMATION AT THE RIGHT TIME (VIRT) AND VALUE OFF THE SHELF (VOTS) A FORMULA FOR NETCENTRIC ENGINEERING.....	244
EXECUTIVE SUMMARY	245
A MARKET-BASED APPROACH TO DEFINING INFORMATION VALUE	246
COMMUNITIES OF INTEREST AND THE INTERNET MODEL	248
NETCENTRIC OPERATIONS PILOT PROCESS	248
A NETCENTRIC OPERATIONS “INCUBATOR”	250
PROPOSED INCUBATION PROJECTS	251
MOBILE COLLABORATIVE NETWORKS	252
"PRIVATE" BUT UNCLAS CROSS COALITION COLLABORATION NETWORK.....	253
VALUED INFORMATION AT THE RIGHT TIME (VIRT)	254
DISTRIBUTED, ADAPTIVE OPERATIONS	256
NETWORK-CENTRIC COMMAND AND CONTROL—A LOGISTICS APPLICATION	256
NETCENTRIC BUSINESS PROCESS:	257
NETCENTRIC PREPARATION OF BUDGET EXHIBITS FOR MULTIPLE BUDGET REVIEWERS:.....	257
15. MANAGED SERVICE GIG, PUBLIC PRIVATE PARTNERSHIP FOR NETCENTRIC TRANSFORMATION	258
EXECUTIVE SUMMARY	259
CHALLENGE/OPPORTUNITY	259
BUSINESS MODEL AND TARGETS	260
E-BUSINESS CONTRACTUAL PARTNERSHIPS	262
CHOOSE CONTRACT PARTNERS CAREFULLY	264
COLLABORATIVE VALIDATION AND VERIFICATION FOR SOA.....	264
CONVERT LEGACY CAPABILITY	265
TAB A: PANEL OF EXPERTS	266
TAB B: EXAMPLES OF E-BUSINESS ECONOMY OF SCALE AND INNOVATION	267
TAB C: LIST OF ACTIONS.....	268
TAB D: VALUABLE INFORMATION AT THE RIGHT TIME (VIRT) ECOSYSTEM.....	270
16. COLLABORATIVE VALIDATION AND VERIFICATION FOR SERVICE ORIENTED ARCHITECTURE	271
17. HARD PROBLEMS IN NETWORK OPERATIONS	276

18. COMMAND RESILIENCY: AN ADAPTIVE RESPONSE STRATEGY FOR COMPLEX INCIDENTS.....	286
19. HASTILY FORMED NETWORKS	288
20. BROCHURE FROM W2COG SYMPOSIUM	295
INITIAL DISTRIBUTION LIST.....	304

BACKGROUND AND OPERATION CONSTRUCT

The Global Information Grid (GIG) represents a fundamental shift by the DoD in information management, communication, and assurance to meeting the timely needs of both the warfighter and the business user. Senior OSD leadership recognized that implementing this vision would require vastly higher levels of collaboration across heretofore autonomous organizations and would need leveraging of commercial technologies augmented to meet DoD's mission-critical user requirements. They also recognized that the current DoD acquisition landscape neither provides incentive to nor convenient processes to encourage cross domain collaboration. They were encouraged by successes in the e-business private sector that have found ways to adopt collaborative practices to achieve competitive advantage in the marketplace.

Accordingly, the Office of Force Transformation, ASD Networks & Information Integration, DUSD Advanced Systems and Concepts, and DARPA, collectively provided \$1.6M “angel money” to the Naval Postgraduate School to establish the World Wide Consortium for the Grid (W2COG) research initiative. The project objective is to create a self-sustained not-for-profit consortium that applies the Internet “open” e-business process to accelerate the GIG. A successful “open” forum is one wherein participants are motivated for their own reasons to make voluntary contributions to rapidly achieve mutually desired “good-enough” approaches to interoperability.

The W2COG research initiative discovered that the principles of netcentric operations (NCO) can be effectively applied to engineering. That is, self-synchronized teams of vendors taking their cue from the Open Source movement can rapidly bundle their separate products to create incrementally more powerful information processing capability. This idea led to two central themes: The first is Valuable Information at the Right Time (VIRT); the power of NCO is in enhancing information utility, not moving data. The second is Value off the Shelf (VOTS); the power of netcentric engineering is in creatively re-using interoperable (i.e. off-the-shelf) components. Hence, W2COG projects demonstrate quantifiable improvements in information processing capability by bundling excellent off-the-shelf components.

The W2COG research initiative spawned the not-for-profit W2COG Institute in July 2005 to establish and maintain the infrastructure required to achieve the goals of the W2COG vision.

The tenets of the W2COG Institute are to 1) create a forum and facility to discover commercial and government best practices and solutions for network and collaboration technologies; 2) establish and maintain a readily searchable data base of government, industry, and academic experts in operational, engineering, and programmatic aspects of net centric operations required to create the solution(s) required; 3) demonstrate the utility of off-the shelf network and collaboration components, and how they can be bundled and used to rapidly satisfy NCO requirements, to include placing them on commercial or government procurement schedules; 4) produce documentation that accompanies the successfully demonstrated bundled capability, and to perform independent testing and validation of the solution (“Underwriters Laboratory” function); and 5) sponsor Research and Development projects and provide grants in areas pertinent

to networking and collaboration technologies and their Independent Validation and Verification.

W2COG Outcomes

The \$1.6M dollars of OSD W2COG seed funds, equivalent to the cost of a typical DoD think tank study, has spawned a functioning global community of experts who have delivered a number of successful and well documented netcentric pilots. The aggregate value of these pilots far exceeds the investment, and proves the hypothesis that the W2COG “open” e-business model can accelerate GIG development. . A partial list of these successes follows:

1. Delivered rapid, concise, cost-effective, impartial and expert analysis of important “net-ready” issues.
 - a. Developed an IPv6 transition strategy document adopted by Australian Defense Force. Project took about six weeks, was conducted asynchronously by government and civilian experts in UK, US, and AU. (See appendix (1))
 - b. Provided focused comment by a team of world-class SOA experts on DISA proposed standards for enterprise services. Project took about two weeks and delivered consensus, concrete, recommendations from broad commercial and academic perspective for achieving DoD’s specific objectives. (See appendix (2))
 - c. Formed a team of world-class experts to conduct a critical, rapid turn around, review of commercial trends and best practices as applied to NORTHCOM requirements for wireless and hastily formed networks. We are awaiting final adjustments to statement of work by NORTHCOM. (See appendix (3))

A “net-ready” analysis document, prepared by a W2COG panel of 4-12 distributed experts, and informed by a global community of hundreds of additional experts, will typically cost a few tens of \$K, and take 30 to 60 days to prepare. Industry studies used to support DoD decisions generally cost \$1-3M and take six to eighteen months to prepare.

2. Created a functioning GIG “data strategy” community of interest (COI). Semantic web technology is critical to the netcentric objectives of the GIG, but is immature. DoD has determined that federated GIG COI semantic data strategies are required, and is developing policy on COI formation and governance. Meanwhile, big software companies are looking for a “killer app” to take semantic web products to market. W2COG has succeeded in selling DoD C2 “semantic track” as that killer app to an impressive COI of IT industry giants, small innovative companies, and government labs. The W2COG Valuable Information at the Right Time (VIRT) COI is applying semantic web technology to model and monitor C2 *track* data among distributed sensor, weapon, and command & control systems to selectively and automatically populate User-Defined Operational Pictures (UDOP). (See appendixs (4) – (7))

VIRT COI members have invested their own R&D resources, conservatively estimated at \$1.5M, to translate the concept into a pragmatic architecture. The team has put

together functioning hardware and software and is ready to apply it to DoD requirements. W2COG is in the process of establishing a similar “Tactical Chat” COI around open source bandwidth-efficient collaborative-service-over-IP tools, methods, and mission threads.

3. Created “netcentric incubator” process to demonstrate, validate, and rapidly field net-ready components for a coalition tactical “edge of the GIG” network.

a. Bundled a suite of COTS web services that allow participants to control multi-level access to collaborative networked “private enclaves.” Information Assurance (IA) is a critical issue in any netcentric activity. First responders to a disaster, or soldiers fighting insurgents, need to share and/or protect information selectively depending on the situation. This reference implementation deliberately demonstrated the capability at an UNCLAS level. It uses NSA-approved commodity-based technology that offers identity management and information assurance similar to that used to conduct on-line Wall Street transactions. This cross domain solution, accredited for SECRET-high if desired, can be reproduced by any government consumer by adhering to the reference implementation documentation and purchasing the hardware and software components from the GSA schedule. (See appendix (8))

b. Partnering with First Marine Expeditionary Force (I MEF) to rapidly field COTS/GOTS sophisticated, covert, wirelessly networked perimeter protection devices for US forces in Iraq. This project links COCOM, government labs, vendors, operators, and MARCORSYSCOM in the same engineering team that will deliver net-ready working product, suitable for DoD-wide off-the-shelf procurement, within six months.

c. Enlisted Prof. Brian Steckler’s NS Humanitarian Assistance/Disaster Response (HA/DR) project to serve as a field lab for W2COG. HA/DR is a principal focus of W2COG because not only is it a vital concern for homeland security, it also offers a high stress UNCLAS test environment and a proxy domain for Stability Operations. The HA/DR response COI of voluntary vendors and government researchers assembles COTS components on the fly to provide “instant internet” in places with no power or communications infrastructure. Team members were on the ground providing internet connectivity and basic C2 capability to first responders in the post-Katrina disaster zone three days after the Hurricane hit. Positive response to Prof Steckler’s work among the W2COG community and others in DoD led to a thematic focus on Hastily formed Networks (HFN) by the NPS Cebrowski Institute and a continuing body of work highly relevant to accelerating fielding of netcentric capability to support, e.g. homeland defense. (See appendices (9) – (10)).

d. Responded to a request for proposal from the Business Transformation Agency (BTA). Mr David Scantling is the senior executive brought in to government from industry to establish BTA’s methodology for certifying and fielding DoD enterprise level software. Mr Scantling had envisioned creating a process similar to the W2COG Netcentric Incubator and was pleased to learn that he can simply apply it to his requirements. We are awaiting BTA review of the proposal at appendix (11).

e. Built a functioning model for a “GIG-lite” on line run-time repository of network-enabling components. These components can be evaluated in context with embedded

mission threads. GIG-lite idea can be scaled to serve as a shared “sand box” for industry and government GIG developers. (See appendix (12))

f. Fielding a reference implementation documentation of a scaleable Netcentric Enterprise Service for multi-level secure PKI identity authentication and performs dynamic role base access authorization, and post event audit. The reference implementation will be developed as part of a PMW160 Limited Objective Experiment last summer 2006 and will be built in partnership with SPAWAR System Center Charleston as an upgrade to capability originally fielded in the Horizontal Fusion project. The technology is described in appendix (13).

g. Proposed a plan to establish a DoD Netcentric Certification Office to work in conjunction with the W2COG Incubator. Function is to provide government validation, verification, certification, and/or accreditation as an embedded activity of the netcentric incubator. The proposal was accepted and funded by DISA Joint Interoperability Test Command. (See appendix (14)).

h. Fielding a reference implementation of a wireless “microserver” netcentric architecture around the aviation industry just-in-time maintenance use case.

Netcentric “incubator” projects will typically spiral in 90 day increments from demonstration to prototype to product. Sponsor seed funds, usually less than \$500K, will leverage partnering vendors’ internal development resources. Successful pilots will deliver government-approved net-ready information processing components immediately to the GSA or other approved procurement vehicle for immediate use by other operational units and/or program managers.

4. Executed a public relations campaign for accelerating GIG development.

a. Conducted a hands-on W2COG formation symposium at George Mason University in May of 2005. Received mandate from ~130 expert participants to create the W2COG Institute and establish an NCO incubator process. Forty nine large and small companies, government labs, universities and individuals joined the Institute to write its charter. (See appendix (20)).

b. Populated “war fighter panels” for multiple NCOIC and AFEI conferences. These panels, wherein uniformed warriors with recent combat experience interact directly with industry, are inevitably hailed as exceptionally valuable. A particularly compelling example was BG Huggins, Chief of Staff of the XVII Airborne Corps, joining one of these panels, net-centrally, live from Iraq.

c. Delivered invited presentations to conferences of the Object Management Group, AFCIA, NDIA, Global Grid Forum, American Institute of Avionics and Astronautics (AIAA), AFEI, CxO Forum, NCOIC, and IPv6 Forum. Published more than a dozen articles and/or interviews. Established a web site that consolidates information about GIG and/or netcentric technical reference, literature, programs, experiments, experts, and activities. Drafted and/or invited, collected, and circulated several expert white papers addressing critical GIG engineering detail. (See appendices (15)-(19)).

Recommendations to OSD

1. Establish a dedicated level of effort of at least one DoD civilian FTE (direct hire, IPA, or continued NPS research) to perform full time outreach to the e-business sector, and to coordinate inherently governmental validation functions of W2COG NCO incubator activities.
2. Encourage DoD engineering activities to employ W2COG Institute as follows:
 - a. Single point of contact for immediate access to a spectrum of broad cross-industry, government, and academia information processing expertise
 - b. Honest broker (e.g. Underwriters' Lab), per status as 501(c)3 tax-exempt charitable/scientific activity, to evaluate, validate, and verify information processing capability
 - c. Risk mitigation and low-cost, rapid turnaround on information processing concepts, pilots, and prototypes
 - d. Means to certify and immediately field successful information processing reference implementations broadly via consumable off-the-shelf model
 - e. Establish "netcentric productivity metrics" per operational COI information processing performance targets.
 - f. Focus GIG semantic data strategy COI activity in general, and leverage VIRT UDOP COI in particular.

APPENDICES

1. ADF IPv6 Transition Strategy



SOLUTIONSGROUP

The Way Ahead

NCW - BATTLESPACE OPERATIONS RESEARCH GROUP

Delivering An Operational Edge Through Collective Thinking, Applied Science & Systems Engineering

AUSTRALIAN DEFENCE ORGANISATION CDR-01 INTERNET PROTOCOL VERSION 6 (IPV6) TRANSITION PLAN

Issue 1 Revision 1.5 (Final)

Date: 29 July 2005

Ball Solutions Group
Level 2, John McEwen House
7 National Circuit, Barton ACT 2600
Postal: PO Box 3276 Manuka ACT 2603
Telephone: (02) 6270 7777 Facsimile: (02) 6273 8125

© Ball Solutions Group, 2005

This document is protected by copyright and the information contained herein is confidential. The document may not be copied and the information herein may not be disclosed except by written permission of and in a manner permitted by the proprietors of Ball Solutions Group Pty Ltd. This statement does not limit the intellectual property rights of the Commonwealth of Australia under PMSS Tasking Directive 928. Permission is hereby granted for Ball Solutions Group Pty Ltd to make unlimited copies of this document for the purposes of the Commonwealth of Australia Department of Defence.

DOCUMENT VERSION HISTORY

Issue	Revision	Date	Authors	Reason for Change
0	0 - 7	18 May 2005 – 9 June 2005	Paul Burns and The Panel	Draft document versions.
1	1.0	30 June 2005	Paul Burns and The Panel	Creation of Issue 1 Document.
1	1.1	14 July 2005	Paul Burns and The Panel	Workshop issues, and Commonwealth comments to Draft Plan incorporated.
1	1.2	21 July 2005	Paul Burns and The Panel	Additional input following Panel review.
1	1.3	25 July 2005	Paul Burns and The Panel	Reorganisation of introduction section.
1	1.4	27 July 2005	Paul Burns and The Panel	Completed the Executive Summary and Conclusions.
1	1.5 Final	29 July 2005	Paul Burns and The Panel	Final version for Work Package 2.

Introduction

Introduction

This Australian Defence Organisation (ADO) Internet Protocol Version 6 (IPv6) Transition Plan (IPv6TP) has been developed by Ball Solutions Group "BSG" in collaboration with a Panel of UK and US subject matter experts "the Panel". The Panel has members from the IPv6 Forum, QinetiQ, the Naval Post Graduate School (NPS) and the World Wide Consortium for the Grid (W2COG). This plan was developed over the period from May through to July 2005 via virtual collaboration (email) between the Panel and three teleconferences between all parties.

A draft version of the plan was delivered to the Commonwealth in June and was the subject of a workshop on 29 June 2005 with the Commonwealth, BSG, NPS and QinetiQ Panel members.

Section 3 of the plan provides the top-down methodology used to generate the recommended IPv6 transition strategy which is detailed in Section 4 of this IPv6TP. The plan provides an IPv6 address space recommendation and includes sections on Governance, Workforce and Risk.

Acknowledgements

BSG would like to extend its special appreciation to Mr Jim Bound (IPv6 Forum), Mr Rex Buddenberg (Naval Post Graduate School) and Mr Chris Gunderson (W2COG) who volunteered their time on behalf of their respective organisations to make crucial and major contributions to the development of this IPv6 Transition Plan for the ADO. The Panel consisted of the following individuals:

Name	Organisation	Title	Role
Paul Burns	BSG	IPv6 Transition Plan Task Manager	Overall task management and point of contact for all personnel and the Commonwealth.
Phil Ashton	BSG	Systems Engineer	Task support
John Pennington	QinetiQ	Senior Principle Consultant - Networks	Contracted to BSG to provide expert IPv6 support.
Jim Bound	IPv6 Forum	Chief Technology Officer	Voluntary provision of expert IPv6 consultancy services.
Rex Buddenberg	Naval Post Graduate School	Professor, Department of Information Science	Voluntary provision of expert IPv6 consultancy services.
Chris Gunderson	W2COG	Executive Director	Voluntary provision of supporting IPv6 consultancy services.

Scope

The scope of this IPv6TP covers the Australian Department of Defence, Defence Information Environment (DIE). Figure 1 indicates that the DIE is composed of Information Domains built upon the Information Infrastructure. Information is currently transported around the fixed and deployed infrastructure by a mix of IPv4 and other non-packetised and/or switched-circuit means. The fixed and deployed infrastructure is composed of an enterprise network and a tactical network.

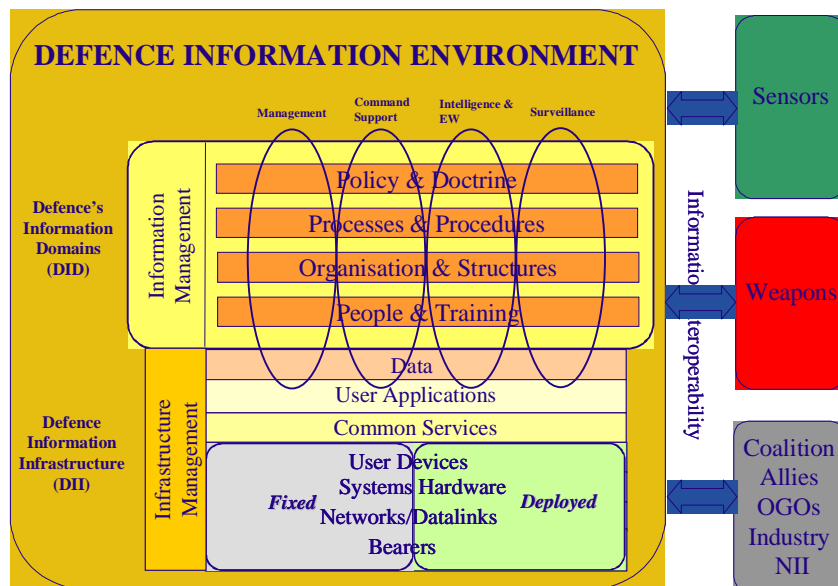


Figure 1 Defence Information Environment

ADO IPv6 Transition Policy

The ADO issued the policy "Transition To Internet Protocol Version 6" [1] in February 2005. The policy states that the transition process will have broad reach across the DIE and involve all¹ Defence computer operating systems, network operating systems, network services, information services, core and distributed networks, and many of Defence's corporate applications.

Policy Statements

The policy states;

- all DIE networks to have completed IPv6 transition by 2013;
- no IPv6 capable hardware or software shall be installed on ADO networks carrying operational traffic unless a risk assessment has been completed, the result approved by the CIOG and use is authorised by the CIOG in consultation with Headquarters Joint Operations Command (J6);
- no IPv6 capable hardware or software shall be installed on ADO networks carrying operational traffic unless a risk assessment has been completed and the result approved;
- the cost of transitioning will be reduced by leveraging information technology (IT) refreshment programs;
- DIE IP enabled hardware and software procured or upgraded that is likely to be in service after 2013 shall be acquired with an IPv6 capability or an upgrade path that will allow it to be upgraded prior to 2013;
- from 1 March 2005 all DIE IP enabled procurements should be both IPv4 and IPv6 capable provided the cost of procurement or the marginal increase in the whole of life cost is acceptable and
- current in-Service ADO equipment that has a scheduled end-of-life before 2010 is exempt from the policy.

The drivers for ADO transition to IPv6 are stated in the policy as follows;

- improved end-to-end network security over IPv4;

¹ It is assumed that "all" relates to all the mentioned systems (e.g. Defence computer operating systems etc) within the DIE.

- better support for the expected growth in the number of mobile IP enabled devices, compared with that provided by IPv4;
- the ability to improve the QoS for IP communications compared with IPv4;
- simpler network management through more efficient hierarchical addressing and routing processes compared with that provided by IPv4;
- is an enabler for the ADO's vision of NCW;
- will aid interoperability with Allies and
- will reduce the likelihood of suffering from technology obsolescence.

The policy also advises that:

- IPv6 migration planning will also develop a consolidated ADO IPv6 address space management strategy to ensure that the ADO's requirements are satisfied and to maximise Allied interoperability and
- the CIOG will manage the transition planning and provide enterprise level guidance on IPv6 transition issues.

This document covers the above "address space management strategy" point in Section 8 and this document as a whole is a part of CIOG's transition planning guidance.

Policy Limitations

The "Transition To Internet Protocol Version 6" policy [1] is aimed at and limited to components of the DIE that;

- are currently IPv4 enabled and will stay IP enabled into the future, or
- components that are not currently IP enabled but will be IP enabled in the future.

The policy does not explicitly mandate² the transition to IP per se of components of the DIE that;

- are currently not IP enabled and are not planned to be made IP enabled in the future³.

Additional Steps To Realising the Policy Objectives

The IPv6 benefits stated in the policy [1] and above in 1.2.1 are not wholly dependent on migration to IPv6, nor will they be guaranteed by migration unless the following significant additional steps are taken:

- Network security
 - High grade IP network encryptors are available now for IPv4 networks. IPv6 capable high-grade products are not yet available off-the-shelf. There is work under way in the US to upgrade the HAIPIE standards to include IPv6.
 - For high grade security there is no significant improvement to be expected from IPv6.
 - The IPSec standards are applicable to both IPv4 and IPv6. Implementations of these standards are widely available, including in Windows 2000 and XP and in most routers. The advantage of IPv6 is that support for IPSec is mandatory and should therefore be provided in all IPv6 capable devices.
 - IPSec VPNs implemented in hosts or routers can be used to provide confidentiality where a lower grade of security is acceptable, for example for 'need-to-know' separation of personnel or financial data.

² Although the policy does not explicitly mandate this, it is recommended that some governance measures should be put in place to ensure that all the required elements of the DIE achieve a routable status in the future to enable NCW, see section 1.2.2.

³ For completeness sake only, there is also the scenario that DIE components that are currently IP enabled could revert "back" to switched circuit (and would therefore not be covered by the policy), however this is not considered as a sensible alternative.

- A public key certificate infrastructure (PKI) is necessary to exploit IPsec. It is expected that the ADO will wish to deploy its own PKI (rather than relying on commercial certification authorities). There may already be a PKI in place to support secure messaging in the ADO, but it would most likely require enhancement to support the more extensive demands of a large IPsec implementation.
- Mobility
 - Mobility is a complicated issue involving potentially many layers of the protocol stack and not just the IP layer. However IPv6 does offer features that can contribute toward improved mobility. If the ADF expects to have increasing movement of users and platforms between networks where the impact is realised at the IP layer, then this feature will be valuable. Please see Annex G for more information on Mobile IP (MIP) and MIP version 6 (MIPv6).
- Quality of Service
 - There is no difference between IPv4 and IPv6 in their support for basic QoS. Although IPv6 packets include a field for flow labels, its use has not been standardized. The QoS field carries the same diffserv code points (DSCP) in IPv4 and IPv6.
 - Effective use of QoS in IPv4 or IPv6 requires ADO-wide agreement on traffic classes, and considerable detailed planning of capacity allocations for each traffic class. If existing service provision contracts do not provide for consistent DSCP definitions, then re-negotiation will be needed.
 - Provision of QoS for Allied networks is an open topic. Currently it is understood that some US networks place Coalition traffic in a different QoS class.
- Network management
 - It is not clear that IPv6 will offer any benefit to network management. There may be potential for some routing efficiencies because of the larger address space, but it will need considerable care in address allocation to achieve this, and the necessary administrative/management overhead may not be justified.
 - The ability of IPv6 to support auto-configuration is often cited as leading to a reduction in management effort. In military secure networks, this must be balanced against the need to have effective control over who or what may connect.
- Address space
 - Although this will be a problem for organizations requiring additional address space, it may not be an immediate problem for the ADO. Unless there are plans to significantly increase the numbers of network elements, then the current allocation should be sufficient. A decision to provide IP capability to all land tactical units would be an example where a significant increase in address space would be required. However, even in this case, a private address range could be used (as the UK MOD is doing within Bowman). A forward-looking long-term view of the potential IPv6 address space requirement is provided in Section 5.
- Interoperability
 - IPv6 everywhere is not essential for interoperability. The extent to which this is required depends on how far the ADF requires network-level interoperability with its Allies.
- Obsolescence
 - Eventually this will be the driver for IPv6 transition. All other issues can be worked around, but at some point it is anticipated that commercial support for IPv4 will be discontinued. The ADO must ensure that all its projects take appropriate action to avoid problems of obsolescence. In most cases this will be dealt with by normal technology refresh activities, although it is noted that refresh

in military systems run over much longer timescales than most commercial IT systems.

Referenced Documents

Publicly Available Documents

This section lists referenced documents including the source if the document is not available through normal Commonwealth channels.

Title	Revision	Date	Source
[1] Defence Information Management Policy Instructions NO 1/2005 : DIE Transition to IPv6		22 Feb 2005	CIOG
[2] IPv6 Essentials, Silvia Hagen, ISBN 0-596-00125-8	1 st Ed.	July 2002	O'Reilly
[3] DoD's IPv6 Transition, Michael Brig, German IPv6 Summit Jun/July 2004		July 2004	http://linda.ipv6.berkom.de/summit/03_mike.brig_DoDs_I_Pv6_Transition.pdf
[4] Function and Performance Specification DWACN JP2047 Phase 2A - Unclassified		31 May 2002	CIOG
[5] Transforming the Defence Information Environment through Improved Governance, AVM Julie Hammer.		5 November 2004	Australian Computer Society.
[6] Internet Protocol Version 6 (IPv6), John P. Stenbit, US DOD		9 June 2003	US DOD.
[7] NATO IPv6 Transition Planning, Rob Goode, NATO Consultation, Command and Control Agency		26 May 2005	Coalition Summit for IPv6, Reston, VA, 26/5/2005.
[8] IEEE 802.16 COTS Technologies As A Compliment To Ship To Objective Manoeuvre (STOM) Communications. R. Guice & R. Munoz.		September 2004	Naval Post Graduate School Thesis.
[9] IP Version 6 Addressing Architecture			ftp://ftp.rfc-editor.org/in-notes/internet-drafts/draft-ietf-ipv6-addr-arch-v4-04.txt
[10] IETF RFC 3587			http://www.ietf.org/rfc/rfc3587.txt?number=3587
[11] IPv6 Response to National Strategy to Secure Cyberspace	Final V2.0	November 14 2002	http://www.nav6tf.org/documents/Response_NAV6TF_Secure_Cyberspace_Final_V2.pdf
[12] NAV6TF PCIPB Input Part II	Final V2.0	December 2 2002	http://www.nav6tf.org/documents/NAV6TF_PCIPB_INPUT_PART_II.pdf
[13] NAV6TF NTIA IPv6 RFC Response	Final	March 1 2004	http://www.nav6tf.org/documents/NAV6TF_Response_NTIA_IPv6_RFC_FINAL.pdf
[14] e-Nations The Internet for All, NAV6TF.		September 23 2003	http://www.nav6tf.org/documents/e-Nations-Internet-for-

Title	Revision	Date	Source
			All.pdf
[15] IPv6 A Practical Technology Maturity Investigation			Naval Post Graduate School Thesis.
[16] Good Network Citizens, Professor Rex Buddenberg.			http://web1.nps.navy.mil/~budden/lecture.notes/good_net_citizen.html
[17] Radio WAN Protocol Notes, Professor Rex Buddenberg.			http://web1.nps.navy.mil/~budden/lecture.notes/r-wan/radio_wan.html
[18] GIG Strategy, Professor Rex Buddenberg.			budden@nps.navy.mil

Government to Government Documents

The following is a list of relevant IPv6 documents that are not publicly available but are expected to be able to be source by the ADO through its Government-to-Government links.

Title	Revision	Date	Source
[19] Navy Internet Protocol Version 6 (IPv6) Technical Transition Strategy	1.0	30 June 2005	US Navy - Karen ODonoghue karen.odonoghue@navy.mil
[20] US DoD IPv6 Transition Plan		March 2004	ADO has this document.
[21] Draft DoD IPv6 Master Test Plan	Rev .13	December 23 2004	US DoD – Michael P. Brig birgm@ncr.dosa.mil
[22] Security Analysis for DoD IPv6 Transition, Report 1: IPsec; NSA report #I333-011R-2004.	Report 1 : IPsec	June 2004	USA – NSA.
[23] US DoD IPv6 Address Plan			US DoD. (This was prepared for submission to ARIN).
[24] Scoping of IPv6 Migration Strategy		July 2004	UK MOD, Integration Authority.
[25] MOD IPv6 Transition Conference, Malvern U.K.		28 February 2005	UK MOD.
[26] IPv6 Issues NC3A, TN-1053	Draft		NATO.

Acronyms and Abbreviations

ADF	Australian Defence Force
AEW&C	Airborne Early Warning and Control
AG	Application Gateway
ATM	Asynchronous Transfer Mode
BSG	Ball Solutions Group
CIDR	Classless Inter-Domain Routing
CIOG	Chief Information Officer Group
CISSO	Command & Intelligence Systems Sustainment Office
DCN	Defence Communications Network
DCP	Defence Capability Plan
DIE	Defence Information Environment
DII	Defence Information Infrastructure
DISA	Defense Information Systems Agency
DMO	Defence Materiel Organisation
DOD	Department of Defence (Australia) Department of Defense (US)
DWACN	Defence Wide Area Communications Network
FISSEO	Fleet Information Systems Support Organisation
FTP	File Transfer Protocol
GIG	Global Information Grid
IPv4	Internet Protocol Version 4
IPv6	Internet Protocol Version 6
IPv6TP	Internet Protocol Version 6 Transition Plan
ISD	Information Systems Division (part of the CIOG)
ISSA	Information Systems Security Assurance (a sub-branch of ISD)
JWICS	Joint World-wide Intelligence Communications System
MAN	Metropolitan Area Network
MANET	Mobile Ad-hoc Network
MOD	Ministry of Defence
NAT	Network Address Translation
NAV6TF	North American IPv6 Task Force
NII	Networks and Information Integration
NIPRNET	Non-secure Internet Protocol Router Network
NTIA	National Telecommunications and Information Administration
OSD	Office of the Secretary of Defense
SIPRNET	Secret Internet Protocol Router Network
TADIL	Tactical Digital Information Links
TIEIO	Tactical Information Environment Integration Office

Executive Summary

This Internet Protocol Version 6 Transition Plan (IPv6TP) has been developed by BSG in collaboration with a Panel ("the Panel") of world-leading IPv6 subject matters experts from the IPv6 Forum, QinetiQ, the Naval Post Graduate School and the World Wide Consortium for the Grid (W2COG). The scope for this IPv6TP includes the whole of the ADO's Defence Information Environment (DIE).

This plan commences in Section 3 by using a Systems Engineering methodology to develop the "context" for Internet Protocol (IP) generally within the DIE and the transition from IPv4 to IPv6 specifically. The results of this "context" setting analysis proposes that "modularisation" is the key to achieving interoperability and "net centricity". Two crucial overall design principles were generated, Principle 1 : Unit Level LANs and Principle 2 : Routable WANs. These principles are used throughout as a basis for many section of this IPv6TP.

The Section 3 analysis also produced derived design requirements for the end-systems that connect to the DIE. The context setting analysis concluded with a definition of the boundary between the non-DIE and the DIE, this is important because the boundary often extends into the ADF's tactical environment and its platforms where many of the "legacy" issues will be encountered in the future.

Section 3 also summarises the IPv6 activities being conducted by the UK MOD, NATO and the US DOD. It was concluded by the Panel that because IPv6 has yet to progress to a sufficient state (anywhere in the world) there are currently no "off-the-shelf" strategies that could be applied to the DIE. As a result of this IPv6TP, the ADO is likely to be in advance of many organisations with regard to its IPv4 to IPv6 transition, and potentially better placed to meet its desired time-schedule if the governance mechanisms can be smoothly and successfully implemented.

The current and future DIE was also analysed with specific emphasis on the DWACN. The future DIE architecture was covered by specifying the DCP projects that will move the DIE from its current baseline to its future state.

Section 3 concluded by providing relevant challenges, opportunities and emerging technologies. The ADO can expect to find its major challenges in the areas of transitioning its non-routable networks and security.

The recommended IPv6 transition strategy is provided in Section 4 and depicted in Figure 15, this shows seven overlapping phases commencing from now until 2013. Importantly this strategy allows for a progressive roll-out of IPv6 whilst recognising that some parts of the DIE may never transition and small enclaves of IPv4 will be required past 2013. The strategy has also been designed to be cost-effective, to have no impact on defence operations and not to degrade interoperability with Allies, justification for this is provided in 4.3.

To reduce the level of risk and ensure a successful transition Section 4.4 proposed a range of information assurance and test activities The recommended strategy section concludes with some specific advice for the key DCP projects.

Section 5 provides a detailed step by step analysis method for constructing a robust IPv6 address plan, this indicates that the IPv6 address range could be anywhere between 34 bits (/30 address) and 46 bits (/18 address). Although this analysis requires further work, it is recommended that the ADO attempt to gain access to the largest contiguous block of addresses possible.

Section 6 details a recommended governance structure for the ADO to transition the entire DIE. Two new organizational offices are proposed to ensure that the governance regime is implemented in an astute and timely fashion and that the actual implementation of IPv6 is appropriately funded and scheduled.

The IPv6 Transition Office (IPv6TO) is proposed to be part of the CIOG, its prime responsibility will be as the “interoperability custodian”. The IPv6TO will become the ADO’s centre of excellence for IPv6 and will also offer technical guidance to the whole of the ADO.

The IPv6 Program Office (IPv6PO) has been proposed to act as the Program Manager for the implementation of IP across the whole DIE. Functionally the office must cover the scope of ADO projects from inception through to second pass (where they are under the control of the CDG) then on past second pass and into service (where they are under the control of the DMO). The IPv6PO is envisaged as an Integrated Product Team (IPT) with members from CDG and the DMO. Its creation, function and lines of reporting are seen as crucial to a successful transition. Section 7 details the organisational structure of the IPv6TO and IPv6PO. Each position within these offices is provided with a position description and details of the required competencies and experience.

The conclusion to the process of developing this IPv6 transition strategy was to assess all its elements (including the proposed governance structure and workforce) for risk, see Section 8.

Transition Strategy Analysis

The section applies a top-down methodology and provides several lead-in and supporting topics (Sections 3.1 to 3.5) in order to generate the “Recommended IPv6 Transition Strategy” which is presented in the following Section 4.

Setting The Context For IPv6

In dealing with narrow topics like how to implement IPv6, it's rather difficult to grapple without a clear and detailed context. Indeed, the justification for IPv6 is weak without this context. It's not clear when or if the ADO will ever run out of IPv4 addresses, therefore the usual address exhaustion reason for justifying IPv6 work is not convincing in the light of real world usage. But placed into an industrialization and network centric context, the case for IPv6 becomes stronger, particularly as a risk mitigation activity.

In order to set an appropriate context we shall use a systems engineering approach, apply a top-down methodology and analyse the following subjects in order:

- Transitioning from artisan-based to industrial based information systems.
- Defining the GIG.
- Over all design principles (These are the principles needed to achieve net-centricity).
- Defining Radio-WAN interface and performance requirements.
- Defining the DIE boundary.

Transitioning from artisan-based to industrial based information systems

A review of the mechanical Industrial Revolution of the 1790s shows us the following:

- A. Use of chemical energy to extend man's muscles (steam and internal combustion engines).
- B. A transition from artisan to industrial methods of building systems. This transition requires an overall design, but then is able to take advantage of specializations of labour to ease constraints on quality and quantity.
 - a. Modularisation of components is essential to the assembly line.
 - b. Standards (e.g. bolt threads) are necessary to the modularisation.
 - c. Technical training is required in the workforce.
- C. Rise of universal public education, where the focus shifted away from Latin and towards maths, chemistry and physics.

These characteristics mirror almost perfectly into the Information Revolution chapter of the Industrial Revolution, this time from the 1990s onwards:

- A. We are using the network and the computer to extend man's mind.
- B. In a muddling way (because we lack historical perspective), we are shifting to a more industrial method of building information systems. This requires an overall design (see the principles below).
 - a. Modularisation is critical to horizontally integrated information systems.
 - b. Standards do not solve the modularisation problem (it's entirely possible to use the correct standards, mis-modularise a system and build a non-functional artisan information system), but the standards are essential to defining the modular boundaries. IP (and IPv6) is one among a handful of critical standards necessary to the modularisation problem.
 - c. We need a technically trained workforce to manage our information systems. The divisions of labour show up on the job survey analyses but not yet in our skill-set definitions and training.
- C. The information technology skill-sets exhibit some patterns that can be capitalised in workforce planning throughout both the commercial and military environments.

Defining The GIG

The US DoD has developed the concept of a Global Information Grid (GIG), however the PowerPoint definitions that are used to describe it are often confusing. For our purposes, the GIG can be defined as the ADO's internet and the definition can be further divided into the following components:

- A. **Terrestrial WAN.** This is analogous to the backbone services provided by the DWACN.
- B. **Unit level LANs.** In the US Navy⁴ there is a good existing proof in the form of the LANs installed their ships. Base/campus area networks are also considered as a good fit into this category.
- C. **Radio-WANs.** The purpose of the radio-WAN is to reach from the terrestrial WAN to mobile platforms (that contain the unit level LANs).
- D. **End-to-end security.** Familiar link and enclave security techniques must be complemented by end-to-end (or object level, or layer 7) security measures as the GIG grows.
- E. **End-to-end management.** It is no longer suitable to manage a network unto itself, the network segment inevitably routes into other network segments and we need to manage end-to-end.
- F. **Upper layer protocols.** The advent of reach to mobile platforms will trigger a new generation of upper layer protocols (MANET, NORM, device-aware, IPv6 and others). This topic remains rather moot until some of the above prerequisites appear (especially the radio-WANs), but is mentioned for completeness. The shortcomings of TCP over satellite networks is a well-known example of current-generation symptoms that need to be addressed in due time.

The GIG is essentially the plumbing for the DIE. What we deliberately leave outside the GIG 'cloud' is the end-systems that attach to it. These end systems (many in mobile platforms) define specific applications. And the end systems also consume IP addresses.

The network plus end systems attached to it can represent information systems (sense, decide, act functions with the communications to connect them together).

Overall Design Principles

This section uses the "industrialisation" observations from 3.1.1 and applies a top-down method to develop a set of core "principles". In this section we focus on how information systems should be assembled, where the aim is to describe a modularisation pattern that all of the ADO's information systems should adhere to.

A modularisation pattern is important to ensure that our systems achieve interoperability across platforms, programs and also with Allies. We can observe that as a side effect of the industrialisation process the life cycles of information systems have become more rational and the cascading maintenance⁵ issues have become much better controlled.

Therefore, we need two principles of 'network centric' to apply to the design and implementation of our information systems.

⁴ In the Blue Navy i.e. the part of the US Navy that sails the open ocean.

⁵ Cascading maintenance: The problem caused when one component (that is tightly coupled) in a system requires maintenance (or replacement) and because of the tight coupling the requirement for maintenance cascades or flows into the other system components.

Principle 1 : Unit Level LANs

Principle 1 : Unit-Level LANs

End-systems (e.g. sensors, weapons, Allies etc) are connected to “the network” and not to each other. They are attached to unit-level LANs which are in turn connected via a router to either a radio-WAN or a terrestrial WAN.

The Unit-Level LANs principle implies that no end systems are connected to each other (e.g. by point-to-point serial links) and no end systems in a platform are connected to off-board entities, that's what the router on the unit level LAN is for, see Figure 2.




Figure 2 Example application of Principle 1

As the implementation of this principle is outside the program scope of the CIOG, it must be captured as a governance issue where the CIOG has directive authority over the ADO's Program Managers to ensure this modularisation and systems view is realised uniformly across the whole DIE (see Section 6).

The end-goal is “modularisation” and it should be kept in mind that “standards”, although important, are only one of the means to that end. Part of the “modularisation” goal is creating what we term “good network citizens”. To become “good network citizens” our end-systems must encompass the following:

- a LAN interface (which implies a protocol stack, which may in turn imply an IPv6 protocol stack);
- an enveloping (wrapper) definition (MIME or XML are good examples), this provides all end-systems (that need to be interoperable) with a common language;
- a means for authentication and encrypting data (e.g. S/MIME);
- setting of DSCP on exiting data-grams for QoS purposes and
- an SNMP agent that affords both local and remote manageability.

Reasonable exceptions

There are some reasonable exceptions to the Unit-Level LAN principle.

The objective is to place the mission sensors, the mission decision support systems and the mission actors (weapons) in an 'inherently interoperable' position. If the platform is, for example,

an aircraft, we should note that this category does not necessarily include the platform's avionics (the information system necessary to fly the aircraft). A mindless enforcement of the above rules on the avionics package yields no interoperability benefits and is likely to be detrimental to issues such as flight safety. How far these rules penetrate into the platform's own control systems should be a decision properly left to the program manager acquiring the platform.

Principle 2 : Routable WANs

Principle 2 : Routable WANs

Make Radio-WANs and terrestrial WANs routable.

The WAN, both radio and terrestrial, can be viewed (SV-1) as a network cloud with routers at the border, see Figure 3.




Figure 3 Routable Network

CIOG has some direct programmatic control over this segment (e.g. DWACN)(which is different to the US CIO organisation), so the issue for this principle is one of self-governance. The next section develops the interoperability requirements for radio-WANs. There are specific requirements (such as for covertness) that are considered to be outside the scope as they are not interoperability related. However if the CIOG is also acting as the Program Manager, then these other specific requirements will also have to be considered.

Defining Radio-WAN interface and performance requirements

The next step is to develop the interoperability and performance requirements for the radio-WANs. Note that it is not necessary to know the specific uses to which these radio-WANs will be put, our approach is to make them part of the general purpose "internet plumbing".

Placing the radio into the rest of the GIG context enormously simplifies the protocol design. By defining a radio-WAN as a network cloud with routers at the edge, we find that we only need to get the protocols in the bottom two layers of the ISO Reference Model correct. Indeed, worrying about layers 3-7 of the Model constitutes an attempt to reinvent the internet (which is clearly a retrograde step) rather than extending the internet to mobile platforms.

There are at least a two ways to analyse the protocol requirements for radio-WANs:

- **Operational views** (Use Cases). This approach is a useful place to start, but it tends to be incomplete.
- **Taxonomic approach**: e.g. look at the IEEE 802 protocols and hypothesize that if a requirement exists here, it's probably something that our military radio-WANs should consider

Operational Views (Use Cases)

Let us consider an infantry soldier as an example. Recalling “Principle 1” (all end-systems attach to a LAN) and applying this to the soldier, he now “wears” the LAN as a part of his uniform. There will be several end-systems that will attach to his LAN:

- his rifle scope (which doubles as a camera and becomes a sensor);
- other cameras (in counterinsurgency operations, it's often useful to snap a picture of a person interrogated and send it somewhere);
- his radio navigation receiver (e.g. GPS) which both tells him where he is and tells his allies where he is (blue force track, or, in USMC-ese, EPLRS);
- his voice communications system (e.g. a VoIP equipment, the microphone and headset of which are part of his helmet) and
- an instant messaging pad (a PDA or something similar that is strapped to his forearm).

For reality check reasons, note that there are voice applications here, but there are also several “data” applications. In order to not burden this soldier with multiple communications systems, the “converged bandwidth” (also known as “all-IP”) solutions are absolutely required.

The infantryman's equipment includes, a router and subscriber station of at least one radio-WAN which plugs into that and becomes the edge of the radio-WAN cloud. Because routers can have multiple ports, this is not mutually exclusive.

Taxonomic⁶ Approach

In applying a taxonomic approach it is useful to dissect the IEEE 802.x protocol architecture. In doing this we are hypothesizing requirements by finding their presence in existing network standards.

Working down from the top of the stack (see Figure 4), all IEEE 802.x protocols (Ethernet, WiFi, WiMAX, etc.) use the IEEE 802.2 Logical Link Control (LLC). This interface definition (known as a SAP – service access point) provides an interface to the “higher layers” in the protocol stack. It is the LLC's presence that makes an 802.x network, a routable one.

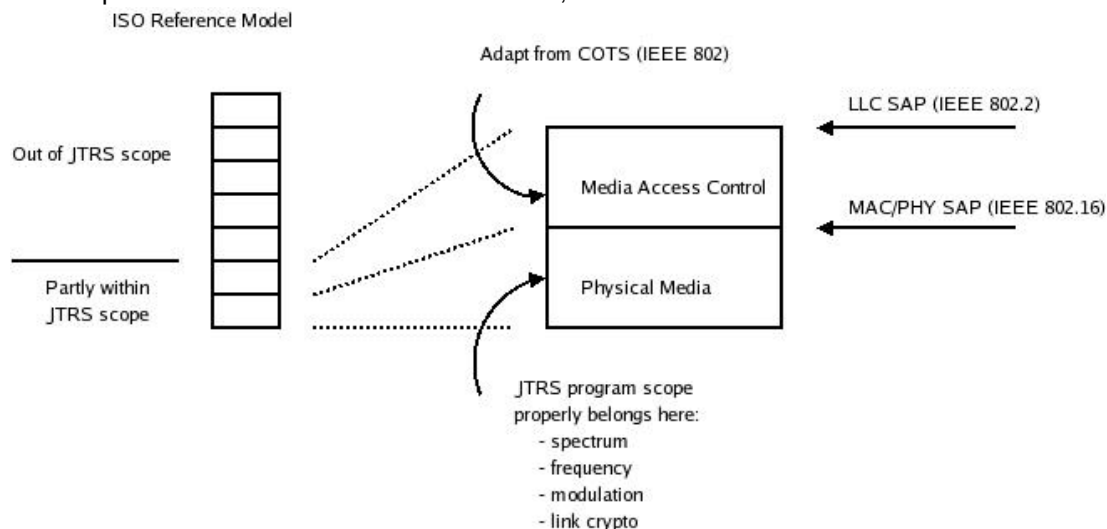


Figure 4 Protocol Architecture

Below the LLC interface is the MAC (Media Access Control) function. The MAC defines how multiple subscriber stations on an 802.x network segment take turns transmitting. There are two kinds of access methods in IEEE 802.x:

- **contention** based access (Ethernet and WiFi both use carrier sense multiple access) or

⁶ Taxonomy : “The science, laws, or principles of classification; systematics.”

- **non-contention** based access (Token systems (including FDDI) and WiMAX (802.16) use this method, these are necessarily more complex, but are more efficient in bandwidth usage and are stable under overload). Non-contention queues also offer ability to control QoS, something that contention based access does not do.

Below the MAC functionality is a security sub-layer. This is a new appearance in IEEE 802.16 and is a reaction to the poorly designed, band-aid approach to security in WiFi (802.11) that turned out to be easily exploitable. The purpose of the 802.16 security sub-layer is to protect the MAC layer messages that pass between subscriber and base stations to control the MAC state machine⁷. The presence of this security sub-layer in 802.16 is an area for further design when considering the additional security measures that should be built into a militarised version of the protocol.

At the bottom of the MAC layer is an interface definition (another SAP) that provides a modular interface to the physical layer (Layer 1) beneath it. This is reflected in some COTS chipsets, particularly for 802.16. Some chip-set vendors have a MAC device and a Physical layer device so the interface has a real-world rendering. This interface is important to us when considering two things:

- adapting COTS technology to military purposes (we want to minimize the parts we have to change) and
- adapting to new technology over life cycles and controlling cascading maintenance (e.g. Ethernet has gone through a half dozen generations in the past 30 years by keeping the MAC stable and changing the Physical layer specification. Even here, the Physical Media Independent (PMI) part of the Physical layer specification has remained stable).

IEEE 802.x splits Layer 1 into two parts:

- **PMI layer.** This is the upper half of the physical layer and contains the frame structures. Essentially all COTS LAN protocols actively used today (e.g. cable modems and DSL) use the Ethernet framing standard. Aside from the COTS reuse aspects, use of Ethernet frames, and the necessary Ethernet addressing scheme, supports multicast. The frame structure is also “protocol independent” meaning that the frame cares not whether the data grams inside the frame are IPv4 or IPv6 or something else.
- **The Physical layer.** This is the lower half of the physical layer and is Physical Medium Dependent. Potential media include wire, glass or the aether. For radio systems (i.e. via the aether), the Physical layer specification includes RF characteristics such as frequency, spectrum, modulation and, if existing, link crypto. Of these, the first three are covered in the IEEE 802.x specifications, usually in exhausting detail.

Management interface

Management interfaces do not map cleanly onto the tightly specified layered protocol stack design despite the fact that IEEE 802.x networks do have management interfaces. In earlier networks (e.g. FDDI) the management interface was captured as a modular specification. In 802.16, the management interface is expressed as a MIB (Management Information Base) within the Simple Network Management Protocol (SNMP) context.

Mapping to radio-WAN programs

There are two major points to consider when mapping to the radio-WAN programs:

1. All of these programs should yield routable networks. This is supported by the GIG context definition, it matches the use case (assuming voice = VoIP) and the presences of the 802.2 LLC in all IEEE 802 networks indicates that this is a solid requirement. We have therefore, necessarily expanded the scope of ‘transitioning to IPv6’ to include the requirement to transition radio-WANs to ones that yield routable networks.

⁷ For readers familiar with current and older generation military satcom systems, these are analogous to orderwire messages.

2. Within the radio-WAN cloud we need to meet four requirements:
 - a. A stable MAC that allows for QoS control. All radio networks can look forward to being saturated, so MAC stability (and consequent bandwidth efficiency) are important. Military networks clearly need to support QoS privileges to some users in some situations (e.g. official business over morale traffic, contact reports over logistics requests, everything else over PowerPoint).
 - b. Multicast. For the "Supply-side", Multicast is the only offset to the limited bandwidth that radio-WANs will have compared to the wired networks they route into. For the "Demand-side", a lot of military data (e.g. the blue force track in the use case) is multicast in nature.
 - c. Layer 2 and 1 security. Clearly we have LPI/LPD (Low Probability of Intercept and Detection), TA/TFA (Traffic Analysis and Traffic Flow Analysis) and jam resistance requirements.
 - d. Management. The ability to manage the components in a radio-WAN is critical, both within the radio-WAN itself and for end-to-end across an internet in which the radio-WAN is a network segment. In the early design stages, its not necessary to derive management requirements via operational concepts, what's important is that all radio-WAN components have SNMP agents embedded in them so that their management interfaces (controls, dials, knobs) can be "read from"/"written to" both locally and remotely in a secure manner.

COTS Re-use

Of the protocols surveyed, IEEE 802.16 offers the best place to start adapting from:

1. Like all other IEEE 802 protocols, it uses the 802.2 LLC so we have a routable network.
2. And most of the objective criteria above is met within the network:
 - a. 802.16 uses a scheduling MAC layer that is highly bandwidth efficient, stable under overload and allows QoS control.
 - b. 802.16 reuses the Ethernet framing protocol and addressing so multicast is easily accommodated.
 - c. The layer 2 security measures in 802.16 are far superior to anything in previous commercial network protocols. 802.16 does not provide layer 1 security – this is one of the adaptations we need to make (neither does any other commercial network spec).
 - d. 802.16 specifies an SNMP MIB.

There are several adaptations required to make 802.16 suitable for use in military information systems, the obvious ones reside in the Physical layer:

- Military users routinely use different spectrum allocations and modulation methods than those used by commercial users. Commercial 802.16 uses higher frequencies and OFDM (Orthogonal Frequency Division Multiplex) modulation. But these changes are confined below the MAC/Physical Layer SAP and do not affect anything above that in the protocol stack (Note: the 802.16 structure could not be used with a HF Physical layer implementation because of the bandwidth requirements).
- Security needs to be added at Layer 1. This can take the form of spread spectrum (which affords covertness in addition to TA/TFA protection). Or it can take the form of link encryption (providing TA). If these protections are provided, we can re-examine whether the incompleteness in the layer 2 (802.16 security sub layer) require further design effort.

Other than the Physical layer there is also one MAC layer problem that needs to be dealt with in adapting 802.16 to a military context. Some of the MAC messages (including some critical ones like the upload map) are transmitted in one frame and are required to be acted upon in the next frame. The existing protocol works (and has been tested by developers to show adequate headroom) as long as frame length exceeds propagation time. In geo-synchronous Satellite Communications situations, the COTS 802.16 protocol would see many frames 'in flight' at any point in time which will cause the timing constraints to be broken. This issue is being studied in

the United States where there is a proposal to solve the problem by simply stretching the MAC frame in time (from the current 0.5 – 2 msec spec in the standard) to whatever the maximum propagation time is.

Use of radio-WANs based on IEEE 802 protocols places the addressing issue beneath layer 3, so the IPv4/v6 questions do not apply.

Managing Legacy

For the purposes of this discussion we put non-routable but current technology (e.g. Link 11 and Link 16) in the same class as legacy technology (i.e. non-routable out-dated technology). TADILs are classed here as legacy because the same methods are used to make them routable as would be used for a pure-legacy system (e.g. Raven CNR). There are two proven methods for handling legacy:

- **Cocooning.** This method uses an 'IP wrapper' around a non-internet communications system. The US Navy's ADNS system employs this method to put IP cocoons around non-IP communications channels such as MILSTAR EHF and SHF channels. There may be cases where the ADO could use this technique, but in the main it is judged to be of limited use.
- **Layer 7 gateways.** These gateways (see Figure 5) provide a means of “entire-protocol-stack” translation from one domain to another. For instance, we can use a layer 7 gateway to translate from the 'pure IP' illustrated above and a platform that has, perhaps, a Link 11 terminal as it's interface to the outside world.




Figure 5 Layer 7 Gateway

The objective is to keep in tact Principle 1 and the “Good Network Citizenship” rules. If either of these is corrupted, all the modularisation benefits will be lost. The means is to add a Layer 7 gateway outside the router, as illustrated. This gateway receives IP data grams with XML-tagged track data from the router. It translates that data into, for example, a Link 11 track transaction and encloses it in a Link 11 frame per all the Link 11 standards. This makes the Link 11 side of the gateway wholly interoperable with a Link 11-equipped platform.

Gateways are not new, nor are they new in this kind of application. But the familiar form may not be immediately recognizable as a gateway. There are a large collection of 'connectors' in Global

Command & Control System (GCCS), including a connector for Link 11. It's a computer-centric implementation of the same tool where this is a network-centric implementation.

Good Network Citizen Data Wrapper

Our definition of a “Good Network Citizen” included the need for an enveloping data wrapper. The list of end-systems must now be expanded to include any layer 7 gateways as well. There is however a severe scalability problem to be avoided. This is because the number of gateways can increase with the square of the number of end-systems to be integrated, e.g. one gateway is required to integrate two end-systems but three gateways are required to integrate (add) a third end-system

The increasingly accepted approach (in the US) seems to be to use XML tagging as the wrapper, if a common wrapper language (e.g. XML) is used, then the exponential effect can be avoided and the number of gateways only expands linearly with the number of end-systems.

DIE Boundary

The DIE does not include the ADF's sensors or weapons but does include the interfaces to allow information to flow between them and the rest of the DIE. Figure 6 illustrates the DIE boundary using the example of a Wedgetail AEW&C aircraft. This shows that the DIE includes the ground to air link and the Link-11 terminal in the aircraft.

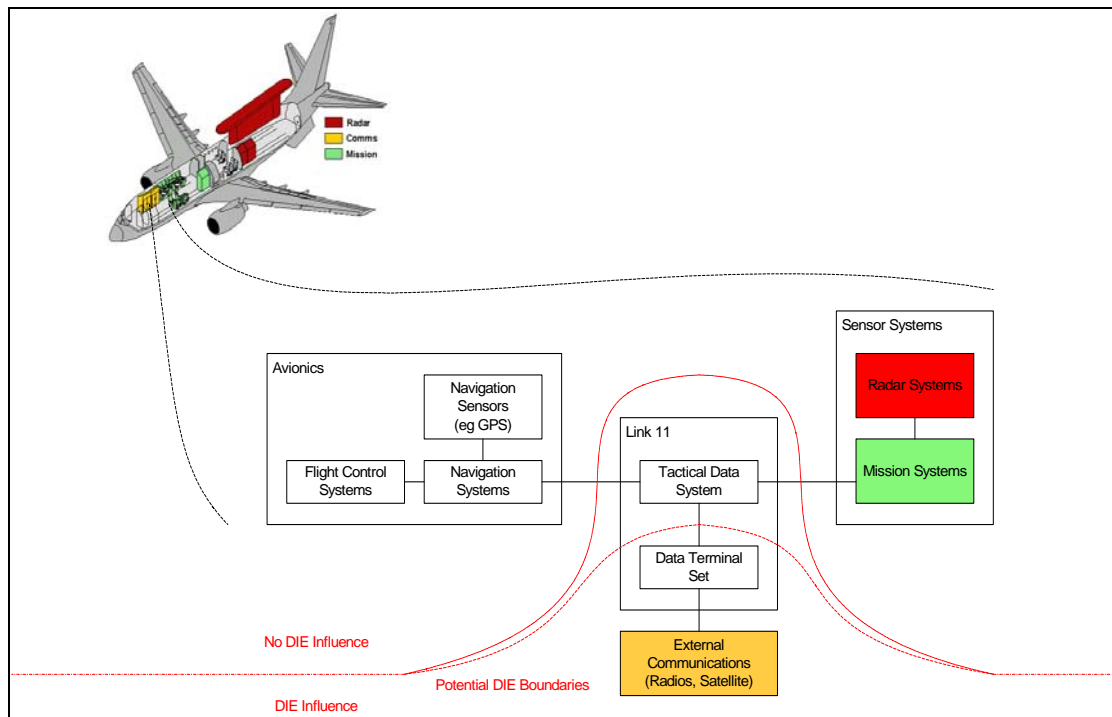


Figure 6 DIE Boundary Example

Figure 7 expands upon the Defence Information Infrastructure (DII) and details examples of both the user applications and the communications systems that make up the static and deployed bearers. The DIE bearers include HF, VHF, UHF and satellite communications systems.

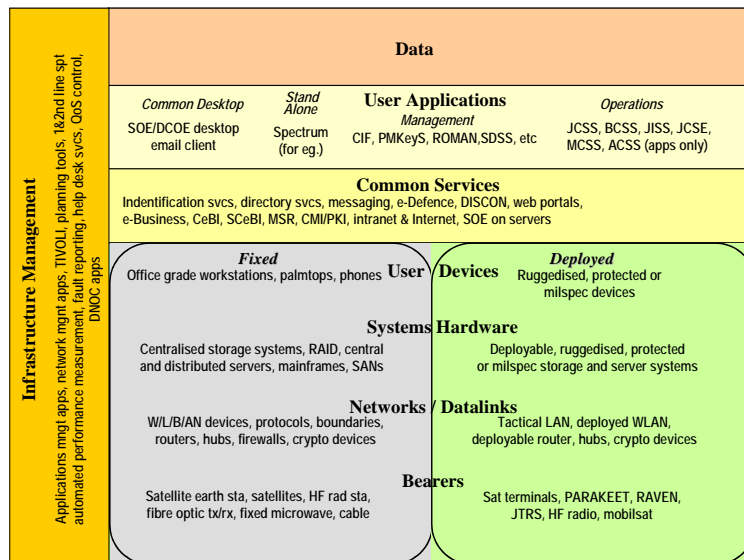


Figure 7 DII (Defence Information Infrastructure) Detailed

Figure 8 shows (an example of) the full extent of the boundary between DIE and non-DIE components of the ADF where Principles 1 and 2 have been followed. In this figure we can see the “mission-thread” from another platform (implementing the decide function) through the radio-WAN to the aircraft.

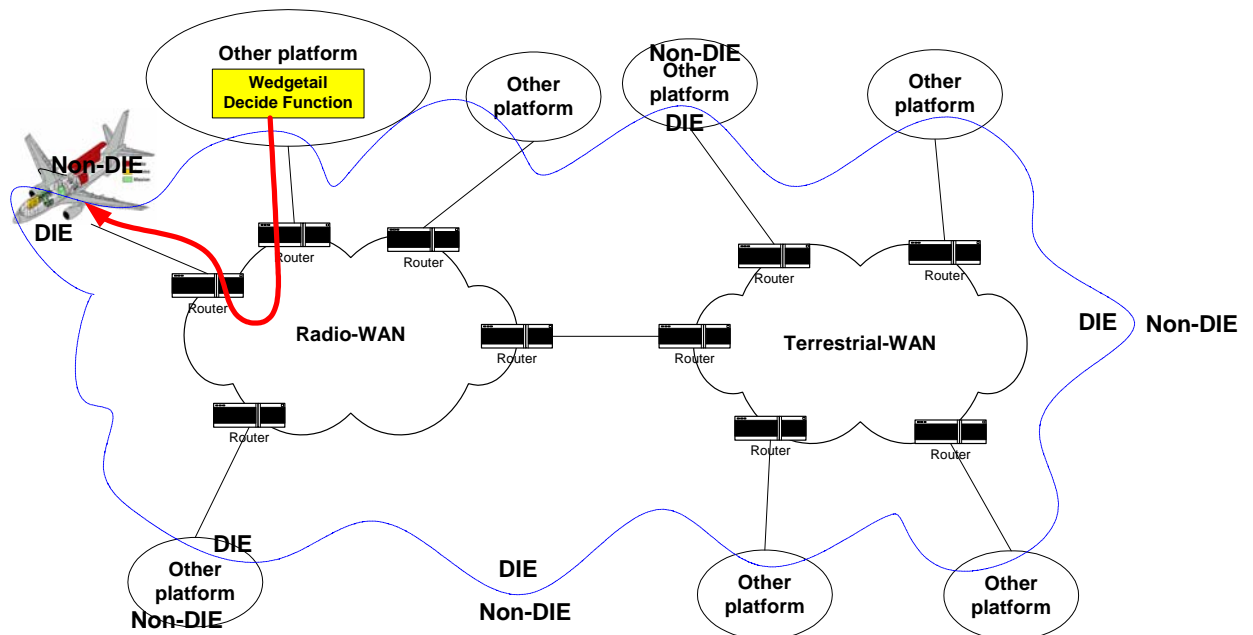


Figure 8 Radio and Terrestrial WANs

Context Conclusion

The above two Principles (3.1.3.1 & 3.1.3.2) will be referred to throughout this IPv6TP.

Execution of these principles modularly separates end-systems and the network infrastructure. This increases modularity, reduces cascading maintenance problems and allows technology insertion⁸ in both the network infrastructure and in the end-systems independent of each other. Implementing these principles does not automatically achieve interoperability, but it does lay the enabling foundation. Conversely, avoiding these principles may lead to interoperability problems within the ADF and between the ADF and its Allies.

IPv6 Background

Current or Previous IPv6 Transitions

The world-wide experience with transitioning to IPv6 from IPv4 is such that the Panel is of the opinion that fully-completed/previous transition strategies do not exist because no organisation has yet completed a transition. The UK, NATO and US defence organisations are in the process of planning their transitions, however the Panel is not able to provide substantial details (beyond what is available in the public domain) of these current transitions strategies because this information either cannot be shared under the existing arrangements and or is classified (See 1.3.2 for a list of Government-to-Government documents).

The Panel has a range of views concerning the actual state of completeness of the MOD and DOD plans. Details of these transition plans can be made available to the ADO by direct liaison between the ADO and the relevant members of these organisations. The Panel will be able to assist the ADO to make the necessary contacts.

The following paragraphs summarise what can be advised concerning the progress with current IPv6 planning within the UK, NATO and US defence organisations.

UK MOD IPv6 Transition

The UK MOD is in the process of developing an IPv6 transition strategy that will be followed by a detailed IPv6 plan⁹. A study undertaken in 2004¹⁰ explored the key drivers for transition and highlighted critical issues. In summary the report concluded that:

- the primary driver for MOD transition is UK – US interoperability;
- there is no pressing UK national need for IPv6 migration;
- the primary UK national driver is to avoid obsolescence;
- the UK MOD has ample address space;
- the features of IPv6 (when compared with IPv4) do not lead to obvious enhancements to military capability and
- security is a critical issue which has yet to be fully explored.

About eighteen (18) months ago the UK MOD set about to quickly determine a strategy for IPv6 transition, and initially settled on a preference to use the Dual-Stack¹¹ approach. That early decision is now being re-appraised and questions have been raised.

In 2004 the MOD Defence Interoperable Network Services Authority (DINSA) was tasked to acquire IPv6 address space. It is our understanding that DINSA do not see this acquisition as a high priority and consequently they have not made significant progress. They are however looking at the request generated by the US DOD with the intention of using it as a template and initially requesting a small allocation with the view to expanding this at a later stage.

To our knowledge there have been no studies to ascertain the address space size or address hierarchy required by the UK MOD. The MOD's current fixed IPv4 network infrastructure is

⁸ IPv6 is one such technology that can be inserted.

⁹ It is understood that the UK MOD and US DOD are considering demonstrating IPv6 at CWID 60/07.

¹⁰ This study was undertaken by a team of two consultants within the Integration Authority from within the MOD's Procurement Agency. One of the consultants was replaced by John Pennington in August 2004.

¹¹ See Annex A for a description of the Dual-Stack approach, its advantages and disadvantages.

provided by British Telecom (BT) who are responsible for technical management of the network including addressing structure.

The UK MOD has yet to allocate funding to conduct an IPv6 study (similar to the one that has generated this Plan/Strategy) and therefore the ADO is likely to be more advanced than the UK MOD in its planning by the time this report is complete.

NATO IPv6 Transition

The NATO Consultation, Command and Control Agency (NC3A) has been tasked to develop an IPv6 transition plan by early 2006. This is aimed at transition for the NATO command chain, which largely operates at the strategic level (between NATO and national defence headquarters). The scope is NATO funded communications and information systems and interfaces to national systems, including national systems deployed in support of NATO operations. NATO is developing a definition of IPv6 conformance and related procurement guidance and a STANAG for IPv6 interoperability may be produced in the future. Initial studies have reached the following conclusions:

- there is no overarching technical reason for NATO to transition to IPv6;
- the drivers are national transition plans and the pace of commercial developments;
- current NATO policy is to prepare for IPv6;
- maintaining operation and interoperability within the complex NATO infrastructure are crucial issues;
- systems will be fully functional and tested prior to cut-over from IPv4;
- a strategy being considered is to cut-over each distributed information system separately and selecting the transition mechanism from a range of options in order to fit the characteristics of the subject system being transitioned;
- the NATO strategy implies that IPv4 and IPv6 networks will be supported in parallel for some considerable time and
- the parallel support of IPv4 and IPv6 will result in increased cost and the chosen transition profile may significantly affect the overall cost, however a detailed cost analysis must be conducted to quantify the cost implications.

This NATO view was reinforced by a presentation at the recent Coalition Summit for IPv6 conference in Reston USA [7].

US DOD IPv6 Transition

The US DOD issued the memorandum "Internet Protocol Version 6 (IPv6)" [6] in 2003. This document provides policy for enterprise-wide deployment of IPv6. IPv6 is specified as the next generation network layer protocol of the Internet as well as the Global Information Grid(GIG)¹², including current networks such as the Non-secure Internet Protocol Router Network (NIPRNET), Secret Internet Protocol Router Network (SIPRNET), Joint World-wide Intelligence Communications System (JWICS), as well as emerging DOD space and tactical communications. The DOD has the goal of completing the transition by the end of the 2008 US financial year. A summary of some of the major elements of the policy include:

- from 1/10/2003 all GIG assets being developed, procured or acquired shall be IPv6 capable;
- transition of the GIG will occur between 2005 and 2007 (US financial years);

¹² US definition: The globally interconnected, end-to-end set of information capabilities, associated processes, and personnel for collecting, processing, storing, disseminating and managing information on demand to war fighters, policy makers, and support personnel. The GIG includes all owned and leased communications and computing systems and services, software (including applications), data, security services, and other associated services necessary to achieve Information Superiority.

- IPv6 was not permitted on networks carrying operational traffic from 2003, subject to further review;
- DISA to acquire IPv6 address space to meet five years of requirement;
- DISA specified as agency to manage DOD IP addresses and
- DOD Chief Information Officer to develop an IPv6 transition plan.

As stated in 3.2.1 the ADO will require direct liaison with the US DOD to share the full extent of its IPv6 planning. The Panel is also aware that the company Electronic Data Systems (EDS) is working on IPv6 for the Navy/Marine Corps Intranet (NMCI) and that the state of this IPv6 planning could be obtained through the appropriate government-to-government links.

The international organisation the IPv6 Forum (www.ipv6forum.org) and the North American IPv6 Task Force (www.nav6tf.org) have provided input to the US DOD to help with their IPv6 transition planning. These two organisations are a significant source of publicly available documentation and work supporting the benefits of transitioning to IPv6.

The paper “IPv6 Response to National Strategy to Secure Cyberspace Final V2.0” [11] highlights many of the problems with today’s Internet architecture (IPv4 and NAT) and is supportive of transitioning from this architecture to IPv6.

The paper “NAV6TF PCIPB Input Part II” [12] agrees with the benefits (put forward in the DIMPI [1]) of adopting IPv6, e.g.

- larger address space for end-to-end global reach ability and internet scalability;
- simplified IPv6 data packet header;
- support for routing and route aggregation, making Internet backbone routing more streamlined and efficient;
- server less (“stateless”) IP auto-configuration, easier network renumbering, and much improved plug and play support¹³;
- security with mandatory implementation of IP Security (IPSec) and
- improved support for IP mobility inherent in IPv6.

The paper also puts forward a (US) business case for the transition to IPv6 and makes a series of recommendations for US Government and US Industry, some of those recommendations included:

- application providers to support Dual IPv4/IPv6 stack (see Annex A for more detail) to begin delivery of IPv6 services coexistent with IPv4;
- take early steps to obtain adequate IPv6 address allocations and
- consider in their (industry) manufacturing plans that the majority of mobile devices, and a growing number of household and consumer-electronic devices will require some form of IP connectivity.

The paper “NAV6TF NTIA IPv6 RFC Response” [13] also supports the previously cited benefits for the transition to IPv6. This paper was generated in response to the NTIA’s request for comment (RFC) on IPv6 and provides a view on the costs of transitioning to IPv6 including “Hardware Costs, Software Costs, Training Costs and Other Costs”.

Although the US DOD IPv6 transition is mandated by [6] to be completed by 2008, it is the Panels view that the actual transition of all IP based systems will take some years longer than 2008. The emphasis for the US is to transition selected IP networks and entities to be IPv6 capable by 2008.

¹³ [12] specifies that, “This is the most important future benefit for the Department of Defense and Home Land Defense communications.”

The Defense Information Systems Agency's (DISA) IPv6 transition strategy is summarised in Figure 9 [3].

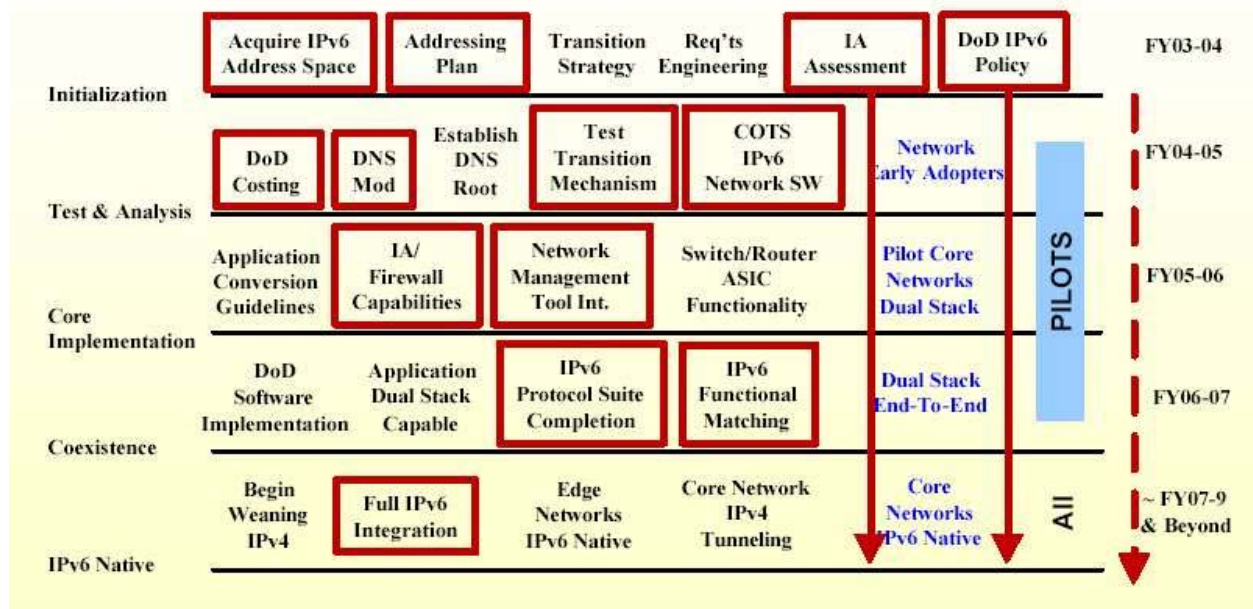


Figure 9 DISA's IPv6 Strategy

DISA does not own any sensors or weapons or decision support nodes (other than their own network operations sites) and therefore the scope of this strategy cannot be considered as a complete US DOD IPv6 strategy. Even if it were completely and successfully executed it would not address the issue of the many artisan¹⁴ data links (e.g. tactical data-links) that are closely coupled to sensors within platforms (Refer to "Principle 1").

Potential Transition Strategies

It is the Panel's view that there are no "off-the-shelf" strategies that could be applied to the DIE, however some elements of the MOD, NATO and US experiences may have potential for incorporation into an effective strategy for the ADO and the DIE. More importantly however the strategy should draw on the principles suggested in 3.1 and work within the timetable and constraints of the DCP programs that will shape the DIE into the future.

The remainder of this transition plan therefore calls upon the collective expertise of the Panel and an analysis of the DIE to provide strategic options and a recommended strategy in Section 4.

Defence Information Environment (DIE)

As an input to forming a suitable transition strategy, it is first necessary to understand the current baseline DIE and then explore where it is likely to progress over the period to 2013. Another view of the DIE (compared with Figure 1) and its interfaces is the view that shows the command support environment, see Figure 10.

¹⁴ Artisan view – Sensors connected directly to decision support, connected directly to actors. Also known as "Stove-Pipes".

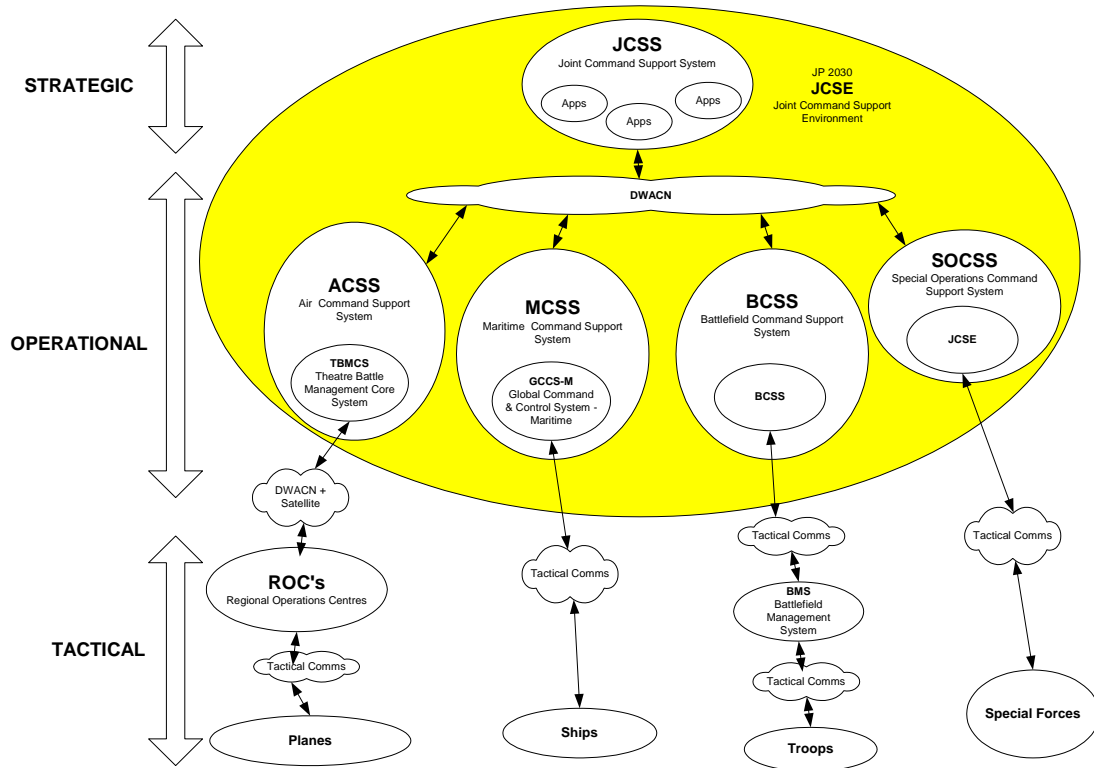


Figure 10 DIE Command Support System Environment

This view indicates that the command support systems are divided along service (Army, Navy, Airforce & Special Operations) lines and each of these systems contains a unique set of legacy DIE infrastructure.

The rest of this section focuses on the current and future configurations of the infrastructure components (DII) of the DIE.

DII Baseline Configuration in 2005

The current (in 2005) DII configuration and architecture description is divided into “fixed” and “tactical” components as follows.

DII Fixed Infrastructure Configuration

A generic view of the Defence Communications Network (DCN) is depicted in Figure 11 and consists of:

- Defence Wide Area Communications Network (DWACN) and
- tactical networks.

The DWACN component of the DCN consists of:

- Defence owned wide area communications equipment/services,
- Telstra owned wide area communications equipment/services,
- Singtel/Optus owned wide area communications equipment/services,
- satellite provider owned wide area communications equipment/services,
- Defence owned local area networks and
- Other Government Organisation local area networks.

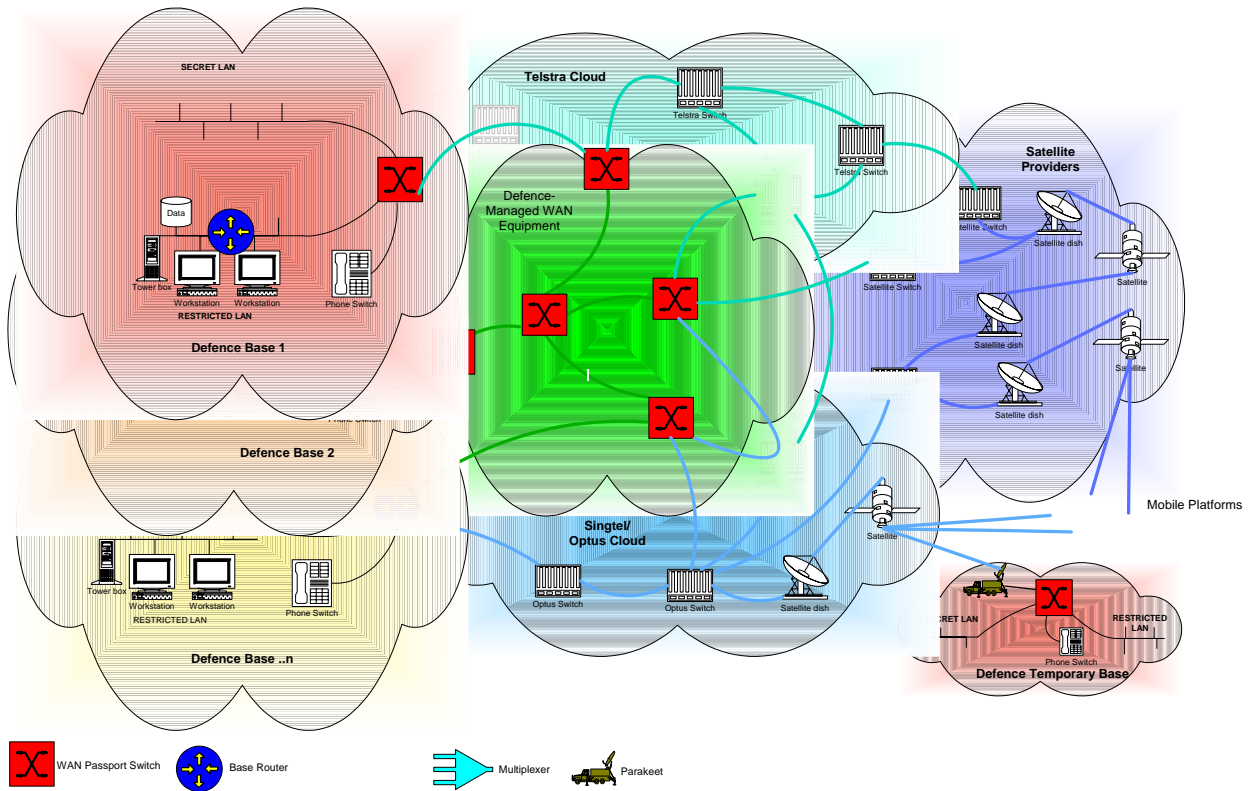


Figure 11 Generic DCN

Defence Wide Area Communications Network (DWACN)

The high level relationships for the DWACN are illustrated in Figure 12. The DWACN provides voice, video and data services via:

- the Defence Restricted Network (DRN),
- the Defence Secret Network (DSN),
- the Defence Voice Network (DVN) and
- other networks.

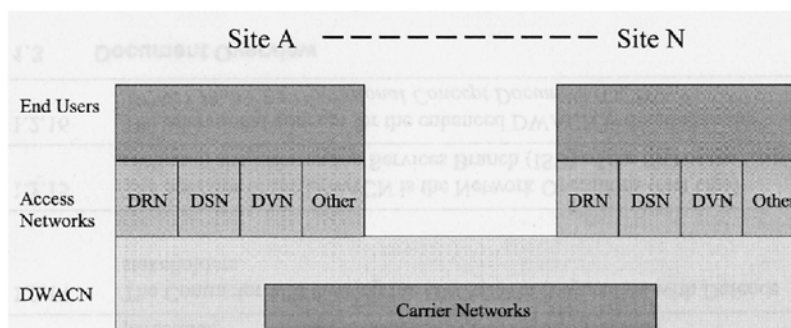


Figure 12 High-level Overview of the existing DWACN Relationships¹⁵

The DWACN interfaces to Carrier Networks (Telstra & Singtel/Optus) and to Defence owned carrier-grade infrastructure. The DWACN provides connectivity between approximately three hundred (300) sites¹⁶, most of which are located within Australia, see Figure 13.

¹⁵ Source = [4] DWACN FPS

¹⁶ Source = [4] 1.2.4 DWACN FPS




Figure 13 High-level External Interfaces¹⁷

DWACN sites include nearly all DOD establishments (bases, barracks, headquarters, offices etc) that are regularly staffed and a small number of Other Government Organisations (OGOs). There are also approximately twenty (20) overseas sites within the DWACN.

The DWACN has a hierarchical structure and is managed centrally. ATM is used extensively to aggregate different types of traffic and traffic from different sources. Commercial telecommunications carrier services (Telstra & Singtel/Optus) are utilised for most of the inter-site transmission.

The core of the DWACN (see Figure 14) consists of fourteen (14) large switches and these in-turn are connected around the core by approximately 150 smaller (Nortel Passport) switches. There is just “one network” and all the DRN, DSN and DVN traffic is handled over this network, all routing and network separation is done virtually by software routers implemented in the switches.

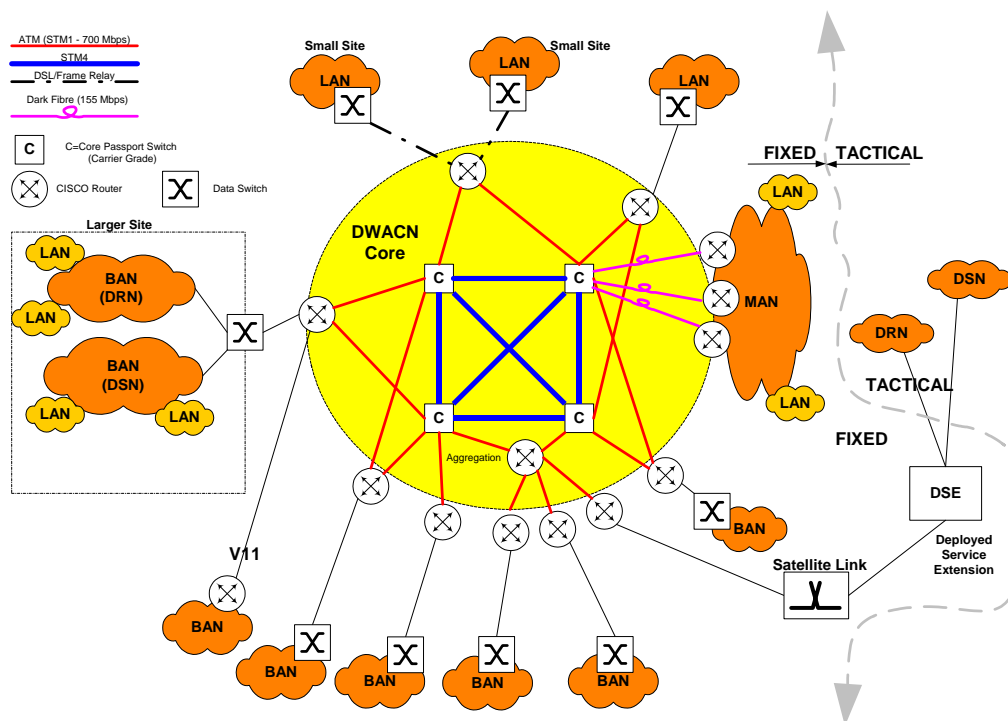


Figure 14 DWACN Core

¹⁷ Source = [4] DWACN FPS

The DRN consists of the following environment¹⁸:

- Users: 79,000,
- Platforms: 55,500,
- Servers: 1300,
- LANs: 500 and
- Applications: 50 Corporate 600 Others.

The DSN consists of the following environment¹⁹:

- Users: 13,000,
- Platforms : 10,500,
- Servers: 395,
- LANs: 70 and
- Applications: 50 Corporate 600 Others.

DII Tactical/Deployable Infrastructure Configuration

The tactical infrastructure consists of:

Airforce Infrastructure

Operational Link 11 in aircraft. Installed but not operational LINK 4A in some aircraft (FA-18) and ROCs.

Army Infrastructure

Raven - Combat Net Radio (CNR) + Parakeet + VHF/UHF/Satellite Trunks.

Navy Infrastructure

Link 11 in some ships (FFGs + FFHs). Extant HF and UHF.

Common Infrastructure

JP2043 HF Modernisation Project. The purpose of the High Frequency (HF) Modernisation Project (JP 2043) is to provide the ADF with a secure, cost-effective information exchange capability for the command and control of deployed forces as a primary survivable system and as a parallel system to satellite communications. The Modernised High Frequency Communications System (MHFCS) comprises a nation-wide network of distributed HF radio stations (the Fixed Network) with a central network management system in Canberra. The Project includes upgrading of HF radio systems in selected mobile platforms and transportable HF communication shelters (the Mobiles). The MHFCS is replacing some of the existing single Service HF fixed-mobile tactical HF gateways.

DII Future Configuration

The future DIE will be shaped by a number of current, ongoing and future projects²⁰.

DII Fixed Configuration

The future configuration of the fixed component of the DII will be formed by the extant components and will be most influenced by the changes implemented by the following DCP projects:

- DEF 7013 Joint Intelligence Support System (JISS),
- JP 2008 Military Satellite Communications ,
- JP 2030 Joint Command Support Environment,

¹⁸ Source = [5] plus information provided by DOD Information Systems Division

¹⁹ Source = [5] plus information provided by DOD Information Systems Division

²⁰ See Annex D for a detailed list of the DCP projects

- JP 2047 Defence Wide-Area Communications Network (DWACN),
- JP 2068 DNOC,
- JP 2069 High Grade Cryptographic Equipment (HGCE) and
- JP 2090 Combined Information Environment.

DII Tactical/Deployable Configuration

The future configuration of the tactical component of the DII will be formed by the extant components and will be most influenced by the changes implemented by the following DCP projects:

Airforce Infrastructure

- AIR 5276 Ph 6 Data links for AP3-C Orion aircraft,
- AIR 6000 Joint Strike Fighter (JSF),
- AIR 7000 Multi-mission Maritime Aircraft (MMA). AP3-C replacement and
- AIR 9000 Helicopters.

Army Infrastructure

- JP 2072 Battlespace Communications System (Land),
- LAND 75 Battlefield Communications Support System (BCSS) and
- LAND 125 Soldier Combat System.

Navy Infrastructure

- SEA 1442 Maritime Communications Modernisation and
- SEA 4000 Airwarfare Destroyer.

Common Infrastructure

- JP 2089 Tactical Information Exchange Domain (TIED) (Data Links).

Challenges

This section identifies the major challenges faced by the ADO in transitioning the entire DIE to a network that will support the concepts of Network Centric Operations. The scope of this section extends beyond simply considering the transition of IPv4 networks to IPv6 ones, but also considers to transition of the entire DIE to IP.

Transitioning Non-Routable Networks

Outside of the DWACN the extant DIE is featured by many non-routable networks²¹. Presuming that there is an aspirational goal (past 2013) to provide routable connections from the network all the way out to the very edge (i.e. to the sensor/shooter), the largest concentration of non-routable entities within the ADF is likely to be in the Army.

Potentially the most significant challenge will be, not only transitioning the extant DIE IPv4 networks to IPv6, but also transitioning the entire DIE toward an all-IP network. Specific challenges are likely to be realised in DIE related DCP programs where:

- they have been in progress for an extended period prior to generation of the ADO's IPv6 policy and a non-routable design has already been chosen;
- there is a DIE component, but the DCP program is ostensibly a platform purchase (Foreign Military Sales (FMS)) where the solution (non-routable) is part of the FMS design or

²¹ Non-routable networks are those networks that do not use IP.

- where the program has sufficient time but insufficient budget to implement a more-costly routable solution. Note that if the design decision is made early enough, the routable solution can be cost-neutral.

For the above cases it is expected that these networks will be very slow to transition to IP and IPv6 with some networks remaining non-routable well past 2013.

The implementation of IP and routable networks is an enabler for interoperability, but does not guarantee it, this leads us to the next challenges.

Interoperability

IP occupies just one layer (layer 3)²² of the seven layer ISO communications model and interoperability between nodes across a network requires all layers at either end of the circuit to be compatible and interoperable.

The transition process will involve the rolling out of hardware and software from various vendors by various organisations at varying points of time. Simply stating that a network component is IPv6 capable will not ensure interoperability with other “IPv6 enabled” components and mechanisms will need to be in-place to test components prior to their insertion into the DIE. However the biggest interoperability challenge will most likely come from differing security implementations cross-ally.

Security

The security mechanisms in place today within the DIE and Allied networks use a mix of physical separation and encryption at the Data Link layer (2) or Network layer (3). Layer 2 solutions have to do with covert communications and link security. The scope of these solutions is limited to single links or network segments where the applied security must be reversed at the nearest router. Layer 3 solutions are those yielded by Virtual Private Network (VPN) equipment as well as firewalls, intrusion detection devices, passwords and physical security measures. Everything within an enclave must run up to the same security level (e.g. Secret). Layer 3 solutions can be vulnerable to insider attacks.

Continuing to solely rely upon these security methods may lead to a future where there is insufficient interoperability within the ADF and between the ADF and Allies to achieve the degree of network centric operations desired. It is not recommended that current security tools be discarded. Link security measures are necessary for covertness and traffic analysis immunity in at least some situations and enclaving (provided by layer 3 security tools) is necessary for infrastructure protection/separation.

It may be possible to improve the flexibility, power and interoperability of the DIE by transitioning to an end-to-end network model with security implemented at the Application layer (7), or “object-based”. This is offered as an “opportunity” and is discussed in 3.5.

Taking this approach and modifying the security architecture would present major challenges (for the ADO and Allies) as it is a fundamentally different approach with potentially large ramifications across the entire DIE. It would need to be very carefully planned by experts and would need to consider both the technical impacts of implementation and the requirements for the development of new policies, practices and training.

Although the concept of object-based security can be recommended for its advantages, it is not placed within scope or on the timeline for the purposes of this IPv6TP. Considering this, the only requirement identified by this plan for IPv6 transition will be to ensure that any cryptographic equipment (implemented at the IP layer) within the security architecture is IPv6 enabled, equipment which only implements security at layer 1 or 2 is not affected.

²² There are other protocols (e.g. MANET ones) that can also occupy layer 3.

Opportunities

This section discusses the lessons learnt and emerging technologies that may influence the transition to IPv6.

Lessons Learnt

Transitioning to IPv6 is yet to have sufficiently progressed anywhere in the world to be in a position to provide any substantial “lessons learnt”. However the Panel has other experience that is worthwhile relating to this IPv6TP:

- Mandating the transition to a new, complicated and competing protocol (e.g. GOSIP) over a short time frame is likely to lead to significant cost with a high probability of failure.
 - Therefore it is recommended that the IPv4 to IPv6 transition is long and overlapping and made as easy as possible.
- The initial “wired” IEEE 80x.x LAN standards (802.3, 803.4 & 802.5) all use a common Layer 2 Logical Link Control (LLC) interface (802.2). As well as all these standards being routable, the use of a common LLC allows all these LAN standards to interoperate and be bridged together. This is also true of the subsequent wireless 802.x standards. The LLC framing standard includes a “payload” field which in the context of this IPv6TP means that any Layer 3 can be interfaced to the Layer 2 and 1 that sits below, in other words these IEEE standards are IPv6 ready. Some of the LAN standards (e.g. 803.4 and 802.5) have faded from popularity but Ethernet (802.3) has continued to evolve with the underlying Layer 1 (PHY) being improved whilst leaving the upper specification virtually unchanged. This success has meant that other standards (e.g. FDDI²³ & DOCSIS²⁴) have reused large amounts of the 802.x standards. Therefore:
 - the features of the 802.x protocol architecture forms a significant basis for the requirements of military wireless networks, with some components being directly applied whilst others will require modification and
 - the 802.x networks are “IP agnostic” and therefore largely immune to the success/failure of narrowly scoped IPv6 transition initiatives.
- The Defence Message System (DMS)²⁵. DMS was conceived in 1988 as a secure messaging system where confidentiality was provided by encrypting parts of the email body. This was designed to provide confidentiality, authenticity, integrity and non-repudiation on an end-to-end, media independent basis. The system took fourteen years (2002) to begin fielding despite the fact that it is essentially a software system. Therefore:
 - there are significant risks in diverging too far from the general trends being followed and developed by the rest of the information technology community and
 - re-inventing an essentially COTS product (email) to provide just one feature (security) has major acceptance risks, on the other hand, adapting essentially COTS products to the same end tends to leverage existing technology and eases user acceptance.
- There are millions of IPv4 nodes in existence today with a very large investment in IPv4 applications, the consequences of this will be that:
 - some IPv4 nodes will never upgrade to IPv6,

²³ FDDI Fibre Distributed Data Interface.

²⁴ DOCSIS Data Over Cable Service Interface.

²⁵ DISA’s description of the DMS. “The Defense Message System (DMS) is the designated messaging system created by the Defense Information Systems Agency (DISA) for the Department of Defense (DOD) and supporting agencies. It is a flexible, commercial-off-the-shelf (COTS) based application providing multimedia messaging and directory services using the underlying Defense Information Infrastructure (DII) network and security services.”

- IPv4 and IPv6 will coexist for an extended period (beyond 2013) that will be heavily dependent on commercial interests and the pace of technological change,
- transition should prevent the isolation of IPv4 nodes and
- it is very unlikely that there will be a “flag day”²⁶.

Emerging Technologies

The following emerging technologies are taking hold within the general Internet and may have some application and advantages for the DIE.

Public Key Infrastructure (PKI)

Public Key technology is used in a variety of places and is becoming extremely important within the Internet. The uses are far more than just securing email and include:

- email security is an excellent example of object level security and as already stated, object-level security is possibly one of the most important enablers for cross-Allied interoperability;
- PKI is used within SNMPv3 (Simple Network Management Protocol). Remote management of a network with any protocol involves exposing data to risks from interceptors and spoofers. PKI enables get, set and trap messages to be authenticated and confidentiality-protected in an end-end, media-independent fashion. This will be increasingly important as networks become more integrated;
- IEEE 802.16 added a security sub-layer in its specification (as a result of the exploits that emerged once IEEE 802.11 WiFi became popular). This layer secures certain MAC-layer messages that pass between the terminal and the base station. The purpose of this security is to increase resistance to man-in-middle attacks which result in theft or denial of service. In the case of 802.16, the PKI is to be managed very similar to the MAC addresses in Ethernet, the X.509 certificates will be factory-installed just like a MAC address and
- various technologies under the general heading of “over-the-air”, such as re-keying cryptographic devices, are using some form of PKI.

IP-over-DWDM

IP over DWDM applies only to the terrestrial WAN part of the large infrastructure. It is not part of the “end-systems”, LANs or radio-WANs which are all on the other side of at least one router. Dense Wavelength Division Multiplexing involves an array of laser diodes at one end of a piece of optic fibre. Each laser is tuned to a different wavelength of light, at the other end of the fibre is found a matching set of photoreceptors, tuned to the same respective wavelength/frequency. This enables very large data-rates to be transmitted down the fibre, 2.4 Gbps and greater.

The great advantage with this technology is that the optic interface can reside inside a core IP router without any intervening technology (other than optic repeaters every ~30 km) between routers, just the fibre optic cable. Switching (frame relay, ATM etc) and ISDN structures (e.g. SONET) are not needed. This leads to a greatly simplified architecture.

Most backbone US ISPs as well as the GIG Bandwidth Expansion (GIG-BE) programs within DISA are using IP-over-DWDM.

Ethernet as a WAN protocol

As IP-over-DWDM was evolving from telephone technology, a parallel development was evolving from the Ethernet community. Particularly with the advent of gigabit Ethernet, the idea of using Ethernet as a long-haul mechanism developed. The basic fibre optic characteristics of Ethernet (e.g. usable distance) are the same as DWDM and the capacities are similar (the industry today

²⁶ A nominated date when IPv4 is turned-off.

has gigabit Ethernet and 10-gigabit Ethernet products that compare well with OC-48 and OC-192 capacities common in IP-over-DWDM implementations.²⁷

Object Based Security

Although it is assumed that the current DIE security architecture will not be radically altered within the period up to 2013, object based security is offered here as an “emerging technology” that may find application in the DIE over the longer term.

By moving the security/encryption function up the protocol stack (from layer 2/3 up to layer 7) to the application layer, it may be possible to improve interoperability between applications. There may also be a performance improvement due a distribution of the encryption function to many terminals.

However a major disadvantage of the “big-cloud” architecture is that the cloud lacks the same type of diversity that exists within the DWACN which currently has several degrees of segmentation that provide diversity and isolate problems. Another issue is that object based security is still immature when applied to austere (radio-WAN) links (<64 kbit/sec) as the issue of validating large certificates has the potential to significantly impact data throughput performance.

IEEE 802.x Based COTS Infrastructure

The IEEE²⁸ has been successful at developing a range of networking standards for wired and wireless physical layer communications. These standards have been turned into successful products (productised) and broadly taken up by the commercial networking community. Successful 802.x standards include 802.3 Ethernet, 802.11x Wireless LAN (WiFi) and 802.16 Wireless MAN (WiMax)²⁹.

One of the keys to their adoption by the networking community has been in their design where the layered communications model³⁰ has been adopted and successful components from earlier standards have been reused in the newer standards, e.g. the 802.2 LLC³¹. This means that the Layer 1 and 2 parts of these standards are payload and Layer 3 agnostic, i.e. IPv6 ready. Some appealing features of these standards include:

- WiMax (802.16) enables routable wireless networks (seamless interconnection to the internet) by virtue of the use of the 802.2 LLC;
- WiFi and more so WiMax, offer wireless broadband at data rates far in excess of those typically in use by the military today³² and
- Large-scale manufacturing, technology advances and commercial adoption have lead to very low cost devices, when compared to military equivalents.

In applying these COTS standards to the military domain the following issues need to be considered:

- range (distance) capability, WiFi's range is purposely limited, WiMax is a better standard here;

²⁷ www.neptune.washington.edu illustrates one program that plans to use gigabit Ethernet as it's (ocean bottom) WAN protocol.

²⁸ IEEE The Institute of Electrical and Electronic Engineers

²⁹ 802.16 is just starting to gain popularity.

³⁰ Ideally each layer (ISO 7 layer model) should only interface one layer up and one layer down, this enables portability between different application at the top and physical transmission mediums at the bottom.

³¹ 802.2 Logical Link Control (LLC) is a Layer 2 protocol (used in 802.3 Ethernet) and can therefore interface to IP (Layer 3).

³² 802.16 has been investigated for use in broadband wireless maritime communications [8].

- WiFi uses a contention based access Media Access Control (MAC) as does Ethernet which becomes unstable under overload and oversubscription and does not allow for QoS mechanisms;
- the requirement is then for a stable MAC that supports QoS³³. WiMax uses a scheduling MAC which provides stability and positive QoS control;
- protocol layer security. WiFi used poor security that was easily broken. WiMax added a security sub-layer (PKI) which provides security for the MAC messages and prevents denial of service and theft of service type attacks and
- physical layer security. None of the commercial wireless standard provide this type of security which is a definite requirement for the military domain (e.g. WiFi uses spread-spectrum which is good for jam-resistance but has a high probability of interception). Requirements such as Low Probability of Intercept/Detection (LPI/D) and techniques including link crypto could be “bolted onto” these standards by replacing/modifying the applicable layer. This is possible because of the adherence to the layered protocol model.
- Timing. Only applies to satellite systems where the (Physical) frame length is exceeded by the return trip propagation time.³⁴
- Multi-cast support.

Recommended IPv6 Transition strategy

This section is the core of the IPv6TP and provides a recommended strategy for the ADO to transition the DIE from IPv4 to IPv6 before 2013.

Strategic Options

Potential options (4.1.1 to 4.1.3) are considered and analysed in the following sections prior to making a final recommendation in 4.2.

Big Bang Transition Option

A big-bag transition would involve the entire DIE being switched from IPv4 to IPv6 almost instantly at some point prior to 2013. This approach would not include a period where IPv4 and IPv6 were run side-by-side. Such an approach is considered to be not only far too risky but it would lead to significantly higher cost than other approaches and is not consistent with the IPv6 DIMPI [1].

Incremental/Phased Transition With Hard Milestones Option

A less risky approach would allow a significant period where IPv4 and IPv6 were allowed to co-exist side-by-side using some or all of the transition/interoperability technologies/mechanisms introduced in Annex A (e.g. Dual-stack, Tunnelling etc).

Implementation of the selected interoperability mechanisms could be mandated “hard milestones” to ensure that the transition is tightly managed and tracked using standard project management techniques. However this approach also does not comply with the DIMPI [1], which specifies that the transition process should leverage technology refresh programs and take advantage of the “natural” progress of “commercial” technology.

Forcing the transition follow a specific timetable could lead to increased cost, interoperability gaps and potentially retard the role-out of IPv6 if commercial technology outpaces any ADO specified milestones. The specified timetable may also limit the ADO’s capability to respond to rapidly changing operational requirements.

³³ 802.16 uses scheduling protocols that meet this requirement.

³⁴ Because some MAC messages are sent in frame x and must be acted upon in frame x+1, otherwise the link is broken and fails.

Incremental/Phased Transition With Soft Milestones Option

The recommended approach is to phase in IPv6 over a long period (between now and 2013) and operate it side-by-side with IPv4. Broad windows “soft-milestones” should be provided to indicate when the various components of the DIE should introduce IPv6 and phase out IPv4, this should be done in accordance with the provisions of the DIMPI [1].

These windows must leverage from technology refresh programs and the planning must be flexible enough to cope with the progress of commercial technology. Currently the core infrastructure within the DWACN (including the Cryptographic equipment) is averaging a five to six year refresh period. With the DWACN being upgraded (JP2047) in the near future it can be expected that there will be up to two hardware refreshes between now and 2013 (see Figure 15). It will also be important to recognise that some IPv4 components will never transition to IPv6 and allowances must be made to continue to support them past 2013.

Recommended Strategy

The recommended strategy is divided into seven phases, these are illustrated in Figure 15 and detailed in the following paragraphs. It is expected that (for DWACN equipment) there will be up to two technology refresh cycles over the period between 2005 and 2013.

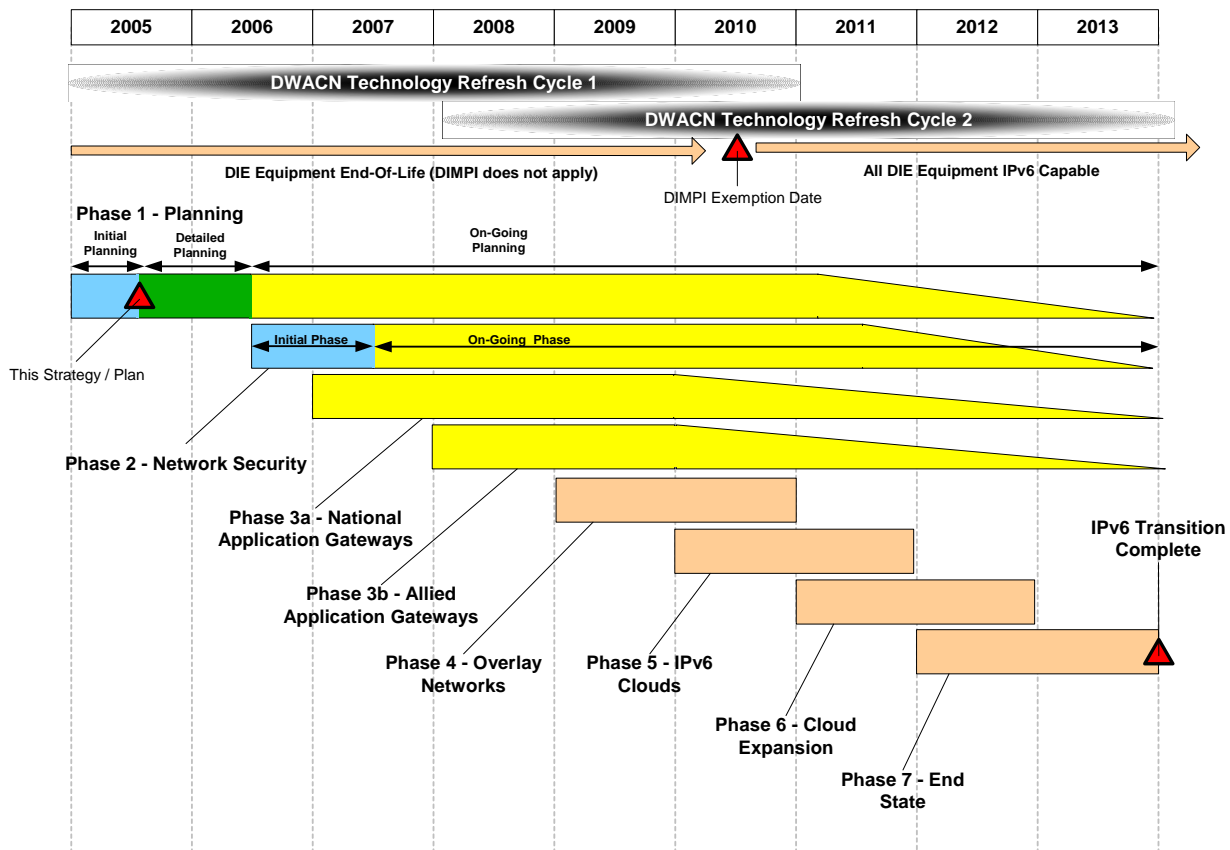


Figure 15 Recommended Strategy Phases for ADO Transition to IPv6

Phase 1 Planning

The planning phase consists of i) initial planning and ii) detailed planning. The ADO IPv6 Transition Plan (this document) forms the foundation of the initial planning phase and provides the big picture view of the whole IPv6 transition process from now until 2013.

A period of detailed planning should then follow the initial phase and this work will seek to answer more detailed questions, see Annex B.

Once information from the detailed planning phase is gathered and understood it will be possible to select the actual IPv6 transition mechanisms and assure network interoperability. This work is likely to be conducted by individual projects but will need high-level coordination by the CIOG (IPv6TO and IPv6PO, see Section 6).

The planning phase will be periodically revisited over the life of the transition to ensure that technology changes are carefully monitored and the state of IPv6 transitions external to the ADO are also considered.

Phase 2 Network - Security

Once the level of detailed planning is sufficiently mature, the Network Security phase will commence. Security will be addressed using two phases:

During the initial phase, security will be enhanced across the DIE to protect against potential threats from the introduction of IPv6. Security will be enhanced by;

- Initial IPv6 threats assessments,
- initially blocking all IPv6 traffic to prevent unauthorised use of IPv6 until protection is adequate) and
- deploying and configuring firewalls, cryptos and intrusion detection systems to provide adequate control of both IPv4 and IPv6 traffic.

This phase is very important to the success of this strategy and should be initiated as soon as possible.

The second phase of Network Security is on going and continues for the life of the DIE. The baseline DIE is continually analysed for vulnerabilities and any threats are treated with counter measures. New network capabilities are thoroughly analysed from a security perspective and only released for use (creating a new DIE baseline) once they are “trusted”.

Phase 3a National Application Gateways

Phase 3a and 3b are not contingent on the previous phase and can commence as soon as enough detailed planning has been completed. National Application-level Gateways are intended to be used intra-DIE, i.e. between disparate DIE networks that need to exchange information. This phase ensures that the various branches of the ADO (see Figure 10) can interoperate at the application level with each other by implementing Application Gateways (AGs) at the network edges. These gateways also decouple the networks, so in principle they also protect against network level threats.

These AG's will allow IPv4 applications (e.g. Email, FTP etc) and IPv4 DIE infrastructure to interchange application level information with like applications on the other side of the AG, independent of the version of IP (4 or 6) being used on the other side of the AG.

These AGs will need to support the range of applications to be used jointly and there will need to be a process of negotiation to determine where the gateways will be hosted and who will be responsible for providing and maintaining them, see 6.3.1.

The AGs could potentially be hosted in any of the ADF support systems e.g. ACSS, MCSS, JCSS (see Figure 10) etc.

Prior to the commencement of this phase of the transition the DIE is completely IPv4. Any infrastructure (routers, switches, servers, hosts etc) that contains an IPv6 capability (e.g. a dual IPv4/IPv6 stack) will have that capability disabled. As AGs are added the matching parts of the DIE can start to transition to IPv6, but as for the Security phase, this phase may need to be run in parallel with the other phases for many years and potentially for-ever if some parts of the DIE or Allied environments (for Phase 3b) never transition to IPv6.

Phase 3b Allied Application Gateways

Allied Application gateways are intended to function in the same way as the National Application Gateways described in Phase 3a, except that they provide a gateway between the DIE and Allied information environments at the application level.

The commencement of Allied Application Gateways is expected to be dependent upon a period of interaction/negotiation with the required Allied IPv6 Transitioning bodies. Because of this, it is recommended that Phase 3a is commenced first followed by an independent Phase 3b that is allowed to run in parallel with the other phases. In this way if the international negotiations take longer than expected, the progress of IPv6 transition within the DIE will suffer no significant impact.

Phase 4 Overlay Networks

The Overlay Networks³⁵ phase can commence in parallel with Phase 3 and begins with the small-scale use of IPv6 applications/systems³⁶ in parallel with a mostly IPv4 DIE, i.e. this phase can commence well prior to 2010. The systems elected to switch to IPv6 are chosen because they need to (or will benefit from) interoperating with other DIE or Allied/Coalition IPv6 systems. Because of the associated coordination issues, it is recommended that the ADO commence with Overlay Networks within the DIE only and then progress to interoperating outside of the DIE with Allies.

For the chosen IPv6 systems, IPv6 data is tunnelled across the DIE's IPv4 infrastructure, through Tunnel-End Points³⁷ and on to the IPv6 end-system, see Figure 17.

There may be some benefit in using "IPv4 compatible IPv6 addresses"³⁸, however this is unlikely to be an effective long-term solution as it may reduce flexibility.

³⁵ These Overlay Networks are intended to be created by tunnelling, which is one way of creating a Virtual Private Network (VPN). VPNs can however be created by other means (e.g. Multi-Protocol Label Switching (MPLS) and this is why we have not used the term VPN. Also, VPNs are sometimes associated with a security function ("private networks") and the Overlay Networks here do not propose any security, just the use of IP tunnels.

³⁶ A small scale IPv6 application/system could consist of anything between one up to several hosts interconnected by a WAN.

³⁷ Functionally either a 4-over-6 or a 6-over-4 tunnel-end point. Note that physically the function usually resides on a router but could also reside on the same machine as a security gateway for instance.

³⁸ "IPv4 compatible IPv6 address is described as "This type of address is used to tunnel IPv6 packets dynamically over an IPv4 routing infrastructure. IPV6 nodes that use this technique are assigned a special IPv6 uni-cast address that carries an IPv4 address in the low-order 32 bits." [2] pg 37.

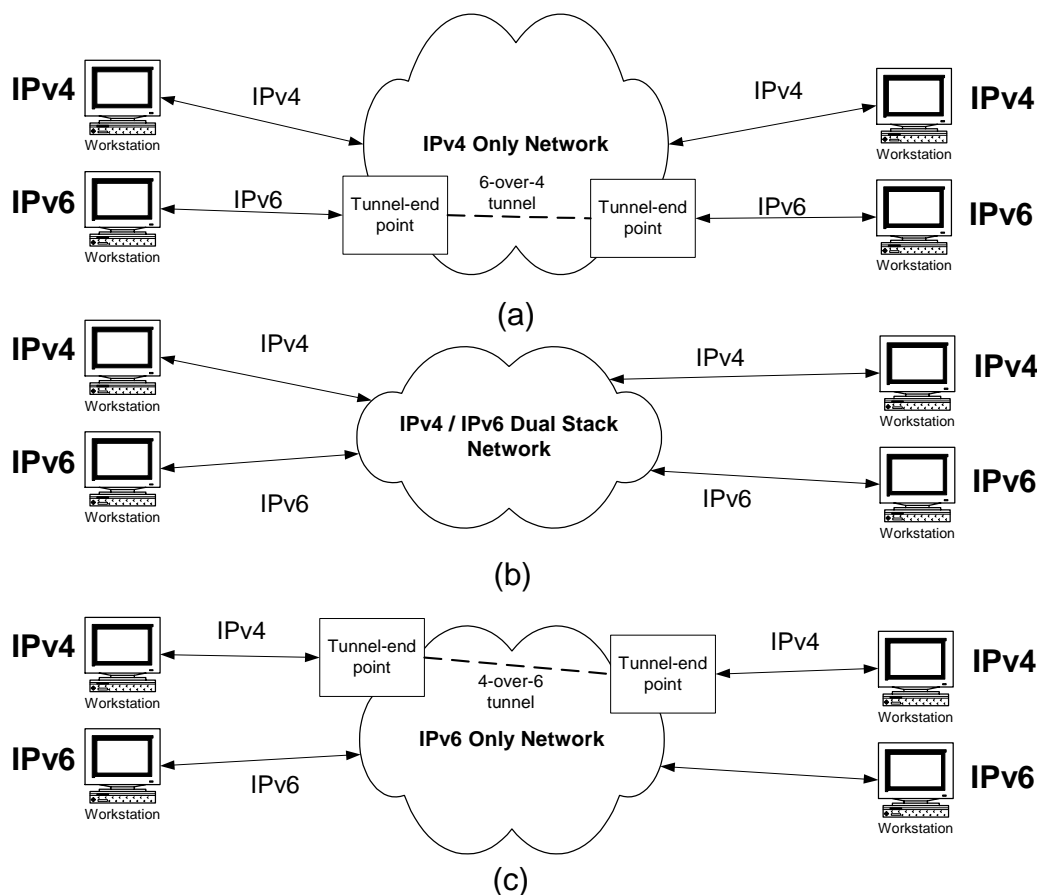


Figure 16 Tunnelling Options

Phase 5 IPv6 Clouds

This phase can overlap with the previous phase and can start whilst the DIE is still migrating to IPv6. This next stage concentrates on migrating larger portions of the DIE to IPv6 along logical boundaries, e.g. complete communication systems, these become the "IPv6 Clouds". Additional AGs and network translation servers are added within the DIE to allow the new IPv6 clouds to inter-work with the rest of the DIE which is still substantially IPv4.

Networks are connected to other networks by using "4 over 6 tunnels"³⁹ and IPv6 clouds are interconnected by using "6 over 4 tunnels"⁴⁰.

Alternatively there may be benefit in taking a dual stack approach⁴¹ (see Figure 16 option (b)) within the IPv6 clouds rather than using tunnelling, especially if bandwidth is an issue or there are security and or fragmentation problems with the tunnelling implementation. The dual stack approach however needs to be analysed for cost (of managing dual stacks) before being considered.

Phase 6 Expanding The Clouds Towards IPv4 Phase Out

The next stage expands the reach of the IPv6 networks within the DIE whilst at the same time shrinking the IPv4 segments, this phase can overlap with the previous phase. This could be achieved by joining together suitable IPv6 systems (implemented in Phase 4) and reducing the number of gateways and tunnels.

³⁹ "4 over 6 tunnels" This assumes that IPv6 only networks are in place and IPv4 packets are required to be sent via the IPv6 infrastructure (See Figure 16 option c).

⁴⁰ "6 over 4 tunnels" (See Figure 16 option a).

⁴¹ Dual-stack is the favoured approach in the UK but may not be necessary or desirable for the DIE.

The expansion process continues until most of the DIE has migrated to IPv6.

Phase 7 2013 End State

This is the 2013 state where ideally all IPv4 systems within the DIE have transitioned to IPv6, however there is likely to be some legacy systems that either cannot be migrated or need to be kept in place because an external party (e.g. Other Government Organisation) is very slow to migrate to IPv6. For this case the required AGs and network overlays will be kept in place as long as required.

Strategy Justification

Cost Effectiveness

The proposed strategy is considered cost-effective as it leverages the natural commercial (COTS infrastructure) refresh cycles that are likely to occur between now and 2013. The strategy does not force any hard-requirements for transition and uses an overlapped phasing plan that will allow for flexibility and co-existence between IPv4 and IPv6 networks and systems.

Impact on Defence Operations

The proposed IPv6 transition strategy is designed to have almost no impact to the ADO at the operational and tactical level. The gradual and phased strategy that allows the co-existence of IPv4 and IPv6 networks should not require any lost capability or “down-time” that is typically associated with large-scale “big-bang” hardware and or software upgrades.

Impact on Interoperability With Allies

An essential feature of the transition strategy is that Allied interoperability will not be degraded during migration, and where possible it should be enhanced. The migration plans of the US DOD will have a major impact in this area, although the ADO will need to co-ordinate with the plans of its Asia-Pacific partners. At present, it is understood that there are few Allied networks. Much of the inter-working is conducted at the application level through appropriate gateways.

For example, the Griffin network currently exchanges information using e-mail with attachments. It would be possible for part of the network to migrate to IPv6, whilst the rest remained on IPv4, with a mail server acting as the interface gateway. This would introduce some additional management cost, and a potential single point of failure. Migration would be simpler if all the Griffin participants agree to transition the network at the same; it should be noted that eventually the e-mail application will need to be transitioned as well as the network.

Other networks in which the ADF participates are CENTRIX and COWAN. These are managed and provided by the US DOD, which will presumably make its plans for transition in consultation with its Allies. If the ADO has application gateways available, it will be possible to maintain interoperability even if the migration timescales are not exactly aligned.

In the future, as the concepts of network-centric warfare are increasingly adopted, there will be increasing requirements for network-to-network interoperability with Allies at the operational and tactical levels. The Allied maritime tactical WAN described in ACP 200 is an example. Migration plans will need to be closely co-ordinated in the appropriate forum. In the maritime case this would be the AUSCANNZUKUS C3I organisation.

Information Assurance and Test Activities

Information Assurance (IA)

It is essential that migration to IPv6 shall not prejudice the security of ADO systems. In this context security includes confidentiality, integrity and availability. It is noted that the US DOD does not yet approve the use of IPv6 networks for operational traffic.

Continuing efforts are required to explore and understand any vulnerabilities which may be introduced by the new or improved features of IPv6. It is recommended that the ADO exploit its close links with appropriate organisations in the US (NSA) and other Allied nations to leverage its national expertise.

IA devices, such as firewalls and intrusion detection systems, must be provided with the capability to handle IPv6 traffic. It is expected that on initial migration to IPv6, including to dual-stack capability, end systems and networks will have some IPv6 features locked down (e.g. neighbour discovery, mobility support). These will only be enabled once appropriate IA protection mechanisms are in place.

Systems migrating to IPv6 (applications, LANs and WANs) will need to be appropriately accredited. It is likely that some systems will only be accredited for IPv6 operation in a stand-alone mode, or only for interconnection over IPv4 networks using secure tunnels. Initially, it is expected that the built-in IPSec features in all IPv6 compliant devices will be used to provide “need to know” separation between communities, rather than military grade security separation. A PKI (public key infrastructure) certificate authority and distribution system will be required to support this.

Military grade IPv6-capable network encryption devices will be required. The ADO may wish to consider taking part in the US-led High Assurance IP Interoperability Specification (HAIPIS) programme.

Test Activities

The ADO will need to gain experience on the behaviour of IPv6 before relying on its use for operational military systems.

It is recommended that the ADO consider taking part in multinational experimental programmes. The CFBLNet (see Figure 17) initiative on IPv6, led by Germany, may be a candidate⁴². This work should focus on IPv4 – IPv6 inter-working mechanisms. The ADO should also initiate a programme to investigate the availability of IPv6 capable network elements and, more importantly, applications.

Initial migration of ADO systems should preferably on a pilot, supporting non-operational information systems, in order to gain confidence.

As systems (applications as well as networks) are migrated to IPv6, they will need to be tested to confirm that the operational requirements are met for performance and inter-working with IPv4. The IPv6 Transition Office (see Section 6.3.1) will oversee this testing, and should be able to reduce the testing requirements as confidence is gained and best practice is shared between projects.

⁴² There are CFBLNet connections at Campbell Park, Russell Offices and DSTL sites in Canberra and Adelaide.

NATO CFBLNet

(for JWID 2004)

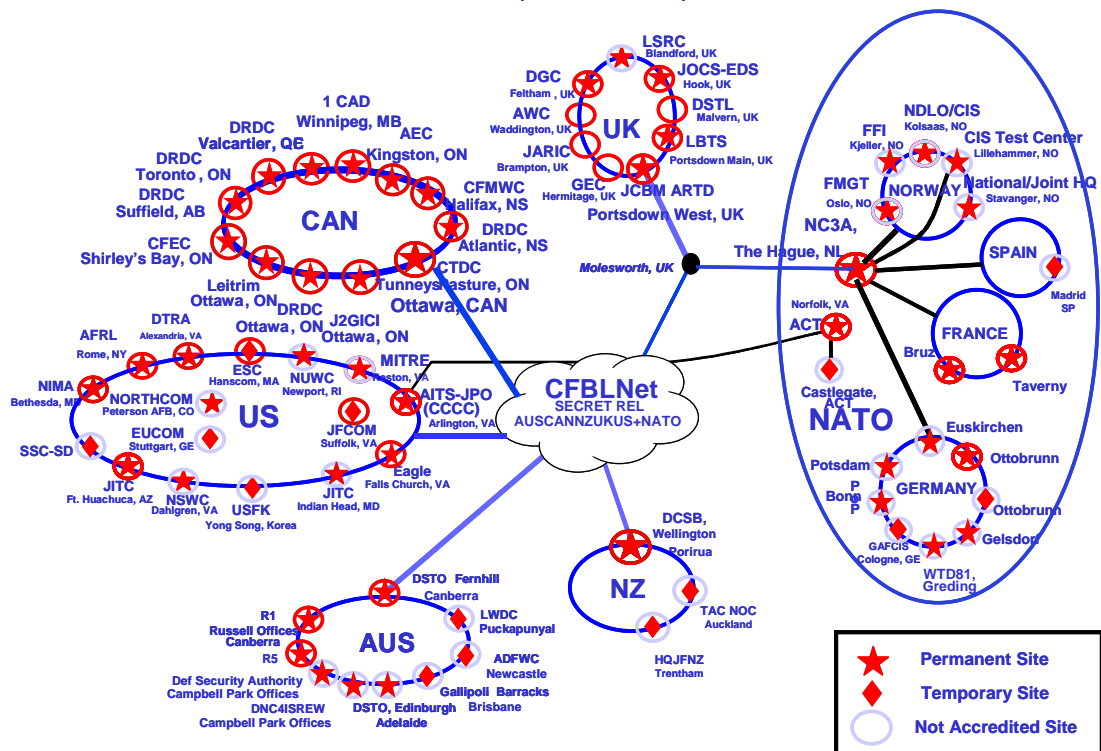


Figure 17 CBFLNet

Key Projects For Transition

A number of current ADO projects will have a key role in implementing the IPv6 transition. This section highlights the actions that these projects should be taking.

- **JP 2047** Defence Wide Area Communications Network

This will be the core programme for transition at the network level. It should develop a strategy and plan for transition including support for IPv6 and IPv4 over an extended period. It is expected that the DWACN plans will drive the planning timelines of other network and application projects. Initial studies should also consider how QoS will be delivered and supported in ADO networks – it is anticipated that “diffserv” will be the underlying technology. This project should also consider how VPN services will be provided, and whether the IPsec features in IPv6 can be effectively exploited. It may well be appropriate for this project to study the provision of mobility services in a general sense across the DWACN area of interest.

The DWACN IPv6 Plan should aim to provide a staged process, where confidence in the reliability and security of the IPv6 service can be gained in a limited environment before extending the scope of the service.

The DWACN currently employs a core ATM switching fabric, based on Nortel passport switches, with Cisco routers at the WAN boundaries. This architecture lends itself well to migration to IPv6, as the underlying connections between core switches use ATM permanent virtual circuits (PVCs). The PVCs are agnostic to the flavour of IP carried over them.

For initial experimentation with IPv6, it is suggested that a small number of dual-stack LANs would be deployed, with edge routers carrying the IPv6 traffic in manually

configured 6 over 4 tunnels. The existing DWACN could provide interconnection between these LANs, using its current IPv4 service.

The second step would be to configure a few DWACN edge routers to provide the 6 over 4 tunnels. Once initial testing is complete, the DWACN could offer a limited IPv6 service to 'early adopter' IPv6 applications. This may also be a useful option for interconnection to allied IPv6 systems.

Step three would be to configure a number of the core Passport switches to support IPv6 as well as IPv4. It might be appropriate to allocate separate PVCs for the IPv6 traffic; this would provide a degree of separation and avoid any inadvertent denial of service to the critical IPv4 traffic.

Once sufficient operational experience has been obtained to provide adequate confidence in the IPv6 service, the fourth step would be to configure all the core switches to support dual stack operation.

A final stage, likely very much later, would be to withdraw IPv4 service and require any legacy IPv4 systems to provide their own tunnels over the DWACN IPv6 service.

- **JP 2008** Military Satellite Communications

If this project includes the provision of services at the network level, then it should develop plans to support IPv6 as well as IPv4. It should be left to this project to determine whether the solution is to be dual stack or tunnelling. The project must also consider QoS support, following the architecture developed in JP 2047. On the other hand, if this project deals with bearer services only, then transition is not an issue.

- **JP 2068** Defence Network Management System and Computer Network Defence (CND)

It is critical that this project can put in place a CND capability for IPv6. It will be necessary to conduct studies on network management during migration. Desirably, a single system should manage both IPv4 and IPv6 network services. Management of tunnels and gateways will also need to be considered.

- **JP 2069** High Grade Cryptographic Equipment

This project will need to ensure that IPv6 capable network cryptos are provided. The capability to pass IPv6 header fields from "red" to "black" is highly desirable, but the security implications will need to be considered.

- **JP 2072** Battlespace Communications System (Land)

For the tactical trunk component of this project, plans should be developed for migration to IPv6. It is probable that these will include support for IPv4 and IPv6, for some period of time, depending on application transition and interoperability issues. The project should be given freedom to determine the preferred approach. QoS must also be provided. The combat net radio (CNR) part of this project will need to give close attention to the IP data capability. IPv6 capable CNR equipment may not be available off-the shelf within the procurement timescale of this project, in which case it will be important to develop plans for migration during a mid-life upgrade.

- **SEA 1442** Maritime Communication and Information Management Architecture Modernisation

It is understood that the maritime tactical WAN will be required to support inter-networking with Allies (ACP 200). Studies on IPv6 migration should take account of the USN's plans for ADNS transition. QoS issues and application transition will need to be addressed.

- **JP 2030** Joint Command Support Environment

It is important to recognise that IPv6 transition impacts applications as much as networks. This project should develop plans for transitioning applications. It should also study how to make use of the QoS capabilities being offered by the networks.

- **LAND 75** Battlefield Command Support System

The same considerations apply to this project as for JP 2030.

Ipv6 address space requirements

Introduction

The aim of this section is to provide the ADO with an analysis method for the DIE in order to make a recommendation for the total IPV6 address space required. The analysis method is designed to ensure that the results enable the expected benefits provided by IPv6.⁴³

The analysis method is demonstrated by providing a worked example (see 5.4), however the results can only be considered as preliminary. It is recommended that the analysis method is revisited as part of the detailed planning phase⁴⁴ (see 4.2.1) and becomes the subject of a specific workshop.

A Case For More Addresses

In direct response to the position that IPv4 address space will meet the world's IP address needs for decades to come, the NAv6TF⁴⁵ has produced the work titled "e-Nations, The Internet for All" [14] (Annex E also provides a view on IPv4 address space exhaustion). This work uses data available from the Regional Internet Registries (RIR) and takes into account the growing adoption of the Internet and networking technologies on a global basis. The NAv6TF view this as a strong and accurate argument for the adoption of IPv6 as the only viable way to sustain the growth of the Internet for all the world's inhabitants.

IPv6 Address Space Analysis Outline

To arrive at an address space plan that meets the long-term need of the ADO, it will be necessary to have a long-term vision for every conceivable network device, node, sensor, and person that may have a requirement for an IPv6 address. It will then be necessary to determine the structure of the network topology that all these addresses will operate from and then interoperate with at other network attachment points. This will determine the prefix size required for the entire ADO IPv6 address space. It is highly recommended that the ADO select an IPv6 prefix large enough to encompass all future addressable network points of attachment.

The IETF IPv6 address architecture document [9] provides the following guidance:

IPv6 addresses are 128-bit identifiers for interfaces and sets of interfaces. There are three types of addresses:

- Uni-cast: An identifier for a single interface. A packet sent to a uni-cast address is delivered to the interface identified by that address.
- Any-cast: An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to an any-cast address is delivered to one of the interfaces identified by that address (the "nearest" one, according to the routing protocols' measure of distance).
- Multicast: An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to a multicast address is delivered to all interfaces identified by that address.

There are no broadcast addresses in IPv6, their function being superseded by multicast addresses. IPv6 addresses of all types are assigned to interfaces, not nodes. An IPv6 unicast address refers to a single interface. Since each interface belongs to a single node, any of that node's interfaces' unicast addresses may be used as an identifier for the node.

All interfaces are required to have at least one link-local unicast address (see Section 2.8 [9] for additional required addresses). A single interface may also have multiple IPv6 addresses of any type (unicast, anycast, and multicast) or scope. Unicast addresses with scope greater than link-scope are not needed for interfaces that are not used as the origin or destination of any IPv6

⁴³ The ADO could also seek a copy of the US DoD IPv6 Address Plan [23], through Government to Government channels.

⁴⁴ To the Panel's knowledge there are currently no publicly available IPv6 address space plans.

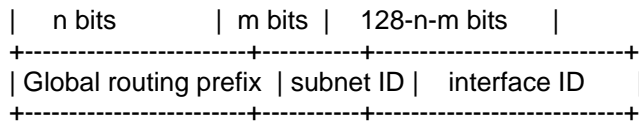
⁴⁵ http://www.nav6tf.org/html/rir_enations.html

packets to or from non-neighbours. This is sometimes convenient for point-to-point interfaces. There is one exception to this addressing model. A unicast address or a set of unicast addresses may be assigned to multiple physical interfaces if the implementation treats the multiple physical interfaces as one interface when presenting it to the Internet layer. This is useful for load sharing over multiple physical interfaces.

Currently IPv6 continues the IPv4 model that a subnet prefix is associated with one link. Multiple subnet prefixes may be assigned to the same link.

The available and current IPv6 Global Unicast Address Format is defined in IETF RFC 3587 [10], and is being used by the Regional Internet Registries (RIRs).

The general format for IPv6 global unicast addresses as defined in "IP Version 6 Addressing Architecture" [10] is as follows:



Where the global routing prefix is a (typically hierarchically-structured) value assigned to a site (a cluster of subnets/links), the subnet ID is an identifier of a subnet within the site, and the interface ID is as defined in Section 2.5.1 of [9]. The global routing prefix is designed to be structured hierarchically by the RIRs and ISPs. The subnet field is designed to be structured hierarchically by site administrators.

After the ADO has determined its IPv6 address requirements and Global routing prefix, it will then need to work with the Asia Pacific RIR (APNIC www.apnic.net). The current work from the IETF regarding Network Address Protection (NAP) [NAP] should also be reviewed. This provides a view of how to define IPv6 networks for privacy and maintain the tenets of end-to-end.

Requirements For Address Space Determination

The important requirements to be met by the address space analysis method include:

- The address space shall be sufficient to permit efficient allocation of addresses to users, equipments and interfaces.
- The allocation process, including registration of names to addresses (populating the DNS) must be fast and easy to manage.
- The address space should be distributed in a hierarchical manner, in accordance with the network topology. This is necessary to facilitate aggregation of routing information, so that the size of the routing advertisements can be minimised. Where networks use limited bandwidth bearers (e.g. long haul links to deployed forces, or tactical nets) this is critically important.
- The address space should be contiguous and therefore the complete allocation will need to be applied for as soon as IPv6 is brought into service.

The requirement to use hierarchical addressing implies that the total address space requirement may be significantly greater (by orders of magnitude) than the total number of addressable interfaces actually used. However, the savings which result from efficient administration and route aggregation will far outweigh the additional cost of address ownership⁴⁶.

⁴⁶ /32 \$2,500 per year (2nd and subsequent years) /20 \$40,000 (2nd and subsequent years)

Analysis Method

Analysis Method Steps

The following is a sequential list of analysis steps that can be followed to design a generic IP network addressing scheme:

- i) Specify the lifetime for the network topology. This is an important step as it will drive the size of the address range required. To ensure that routing efficiency is maximised it is recommended that a single contiguous address space is sought and utilised. It is assumed that the size of the network will only continue to expand from the present day through its lifetime.
- ii) Design a hierarchical network topology to meet the specified lifetime. Start with the network core (parent/top level of hierarchy), then add child/lower level sub-nets in accordance with the operational, security, physical and legacy systems requirements and constraints. Continue adding subnets in a hierarchical manner until the lowest "IP" addressable entity is reached at the end of each of the network's branches.
- iii) At each level in the hierarchy specify the maximum number of interfaces. This is achieved by analysing each branch on the subject level and using the branch that yields the largest number of interfaces. The final result is achieved by rounding up this number to the nearest power of two. This will determine the number of address bits required (e.g. 2 bits = maximum of 4 interfaces, 4 bits = maximum of 16 interfaces etc). In general it is expected that the number of interfaces will increase as one moves down each level of the hierarchy (i.e. Level 2 may have more interfaces than Level 1 and Level 3 may have more than Level 2 and so on).
- iv) Review the address prefix structure and size. If the size is too large then re-visit the levels of the hierarchy where the allocated binary address size is just larger than a binary increment (e.g. 17 is just beyond 16) and review the assumptions to see if one or more bits can be saved by slightly reducing the allocated size to below the previous binary increment.

Worked Examples

In each example, we have followed the usual practice of allocating 64 bits to the interface ID. Typically the interface unique MAC address is used, which allows stateless auto-configuration⁴⁷ to be used, if permitted by the security policy.

Using the analysis steps provided in 5.4.1 we provide the following worked examples for reference.

A Large Network Example

- i) Lifetime is specified as 15 years, this assumes that the design meets the 2020 needs of the ADO.
- ii) The DWACN is assumed to be at the core (Level 1) in the highest level of the hierarchy (see Figure 18). The subsequent levels are populated as follows:
 - a. The next level (Level 2) is occupied by the virtual (uses the same core infrastructure) and other physical security domains. The virtual domains include the DRN, DSN and DVN⁴⁸, the other security domains could include multiple coalition domains, other Australian government domains and the Internet etc.
 - b. The next level (Level 3) is occupied by a number of Base Area Networks (BANs) and a number of Long-haul sub-nets to meet operational requirements. The number of Long-haul subnets will be determined by assessing the number of geographic areas required to be covered and the actual coverage of the available

⁴⁷ RFC2462

⁴⁸ It is recommended that the use of addresses to provide security separation be expressed differently when making the application to APNIC.

Long-haul service options. The Long-haul subnets could be IP transit services provided by ADO, allied or commercial networks.

- c. The next level (Level 4) is occupied by a number of LANs connected to their parent BANs and a number of Tactical area sub-nets connected to the parent Long-haul subnet. The Tactical area sub-nets represent IP service provided over a system such as Parakeet (JP2072 in the future). These sub-nets provide service to a number of deployed headquarters (HQs) as well as transit service to mobile sub-nets.
- d. The next level (Level 5) is only utilised on the Long-haul branches and is populated by a number of mobile sub-nets. The mobile sub-nets could comprise a number of routers installed in land vehicles (Australian Light Armoured Vehicles (ASLAVs), Command vehicles, Jeeps etc). In the future a group of aircraft, or a swarm of Unmanned Combat Air Vehicles (UCAVs) might also form a mobile sub-net.
- e. The last level (Level 6) is also only utilised on the Long-haul branches and is populated by a number of LANs within each vehicle.

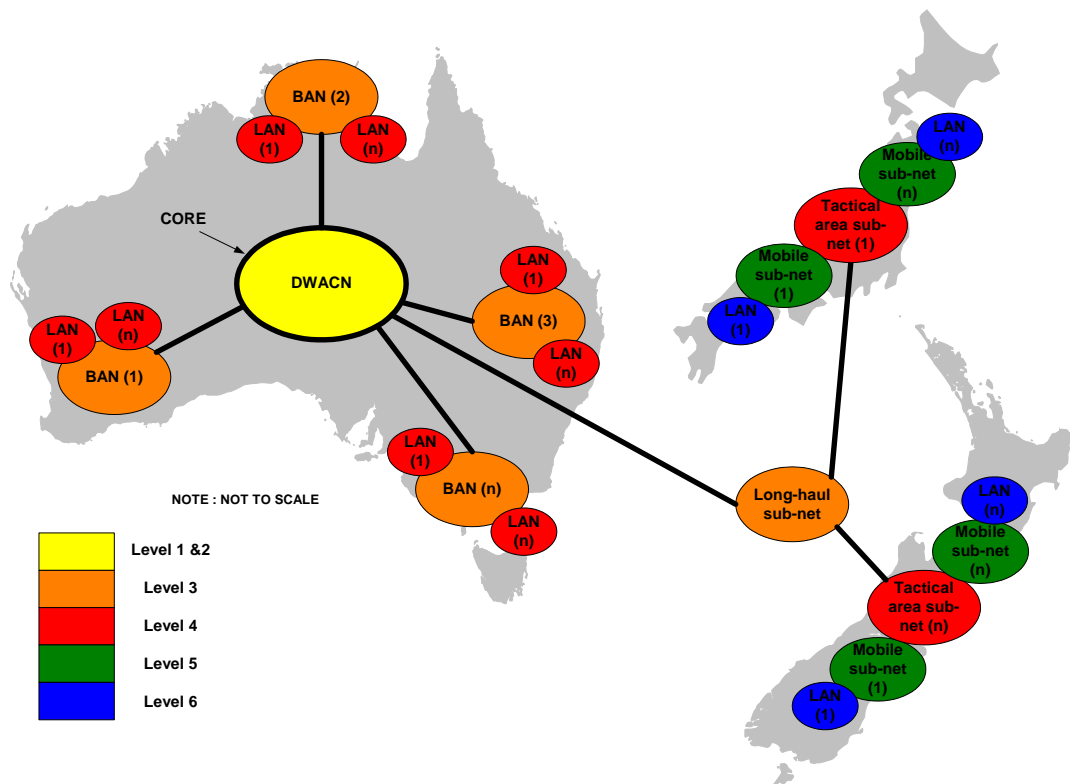


Figure 18 Example⁴⁹ Network Topology

⁴⁹ Except for Australia, these countries are only an example.

- iii) The number of interfaces at each level of the hierarchy is depicted in Figure 19 and detailed as follows:

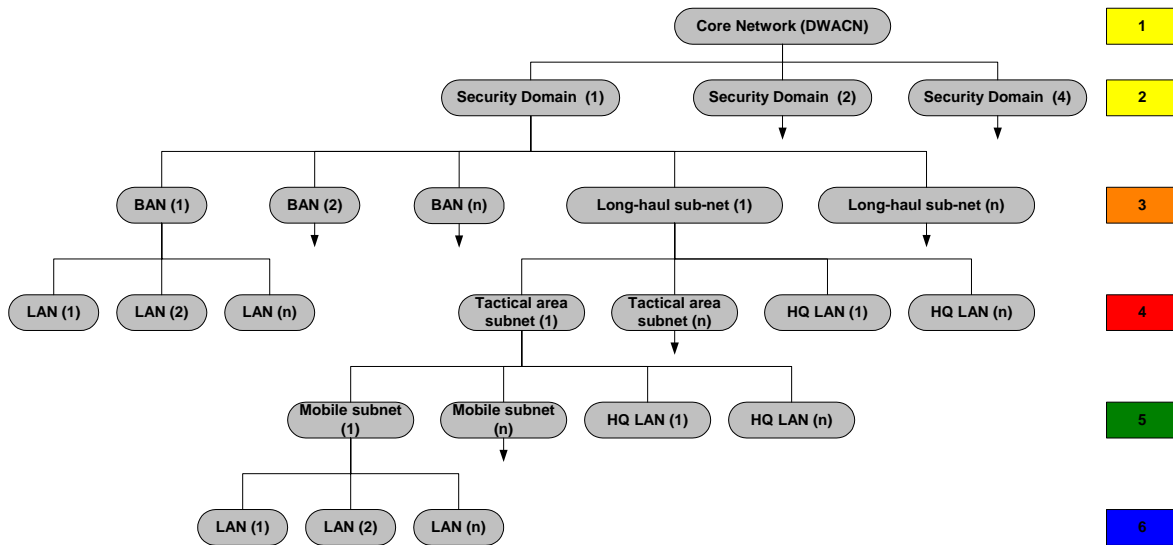


Figure 19 Example Network Hierarchy

- Level 1 interfaces. Assuming that there are up to 4 ADO security domains (including DSN, DRN and DVN), 3 coalition domains, 2 other government domains and 2 miscellaneous domains, this equals 11 interfaces, rounding up to nearest power of two equates to 16 interfaces (4 bits).
- Level 2 interfaces. Assuming that there are 300 BANs⁵⁰ and 4 Long-haul sub-nets and 600 internal routers, this equals 904 interfaces, rounding up this equates to 1024 interfaces (10 bits).
- Level 3 interfaces. For the BAN branches, it is assumed there would be a maximum of 4 LANs and 10 internal routers, for the Tactical sub-net branches, it is assumed that there would be a maximum of 4 Tactical sub-nets, 4 attached Headquarters and 20 internal routers. Therefore the Tactical sub-net branch has the largest number of interfaces (28), rounding up this equates to 32 interfaces (5 bits).
- Level 4 interfaces. As both the (BAN/LAN) and the (Long-haul sub-net/HQ LAN) branches have terminated we only need consider the attachments to the Tactical area sub-net branches. For these branches we assume that there will be a maximum of 4 mobile sub-nets, 40 HQ LANs and 50 trunk routers, therefore 94 interfaces and rounding up this equates to 128 interfaces (7 bits).
- Level 5 interfaces. As the (Tactical area sub-net/HQ LAN) branches have terminated we only need consider the attachments to the Mobile sub-net branches. For these branches we assume that there will be a maximum of 100 vehicle LANs, therefore 100⁵¹ interfaces and rounding up this equates to 128 interfaces (7 bits).
- Level 6 interfaces. At each vehicle LAN we assume that the maximum number of interfaces is 100 and rounding up this equates to 128 interfaces (7 bits).

⁵⁰ Source = [4] 1.2.4 DWACN FPS

⁵¹ The figure of 100, could well be argued and is very much a forward looking number assuming that in 15 to 20 years time there could be many entities requiring and IP address.

- iv) The address structure uses 40 bits (a /24 address) as follows:




Figure 20 Address Size

A Future Large Network

We now consider expanding the previous example by considering potential areas of growth.

- i) Lifetime is also specified as 15 years.
- ii) The same topology and hierarchy is assumed.
- iii) The same number of interfaces at each level is assumed except for those at Level 6. It is likely that major increases in demand from address space will only arise from significant changes in technology. One potential for the additional of an address-hungry sub-network could come from the addition of unattended sensors. These sensors (small low-power seismic, acoustic, RF sensing etc) could be scattered from the air in their thousands across a tactical area. This could lead to a sub-net at Level 6 with say 5000 nodes/interfaces. Assuming that the sensor control station branches from a mobile sub-net this would be rounded up to a maximum of 9182 interfaces (13 bits).
- iv) The address structure uses 46 bits (a /18 address) as follows:




Figure 21 Address Size

A Modest Network

We now consider the previous “Large Network” and revisit each level in the hierarchy to investigate how the network could be reduced in size to a network with a more modest address space requirement.

- i) Lifetime is also specified as 15 years, this is viewed as the minimum requirement.
- ii) The number of interfaces at each level of the hierarchy is as follows:
 - a. Level 1 interfaces. Assuming that there are just 3 ADO security domains (the DSN, DRN and DVN), 2 coalition domains (Restricted and Secret), 1 other government domain and 2 miscellaneous domains, this equals 8 interfaces, rounding up to nearest power of two equates to 8 interfaces (3 bits).
 - b. Level 2 interfaces. Assuming that there are still 300 BANs⁵² and just 2 Long-haul sub-nets and 200 internal routers, this equals 502 interfaces, rounding up this equates to 512 interfaces (9 bits).

⁵² Source = [4] 1.2.4 DWACN FPS

- c. Level 3 interfaces. For the BAN branches, it is assumed there would be a maximum of 4 LANs and 10 internal routers, for the Tactical sub-net branches, it is assumed that there would be a maximum of 2 Tactical sub-nets, 2 attached Headquarters and 12 internal routers. Therefore the Tactical sub-net branch has the largest number of interfaces (16), rounding up this equates to 16 interfaces (4 bits).
- d. Level 4 interfaces. As both the (BAN/LAN) and the (Long-haul sub-net/HQ LAN) branches have terminated we only need consider the attachments to the Tactical area sub-net branches. For these branches we assume that there will be a maximum of 2 mobile sub-nets, 20 HQ LANs and 30 trunk routers, therefore 52 interfaces and rounding up this equates to 64 interfaces (6 bits).
- e. Level 5 interfaces. As the (Tactical area sub-net/HQ LAN) branches have terminated we only need consider the attachments to the Mobile sub-net branches. For these branches we assume that there will be a maximum of just 64 vehicle LANs, therefore 64 interfaces and rounding up this equates to 64 interfaces (6 bits).
- f. Level 6 interfaces. At each vehicle LAN we assume that the maximum number of interfaces is just 64 and rounding up this equates to 64 interfaces (6 bits).
- v) The address structure uses 34 bits (a /30 address) as follows:




Figure 22 Address Size

Regional IPv6 Addressing

IPv6 address allocation is managed on a regional basis (by ARIN, RIPE, APNIC). It is the aim that global routing will be more efficient by allocating address blocks in relation to the location of the organisations requesting the allocation.

The ADO has a fixed infrastructure, but also expects to be engaged in operations with a regional or occasionally global reach. Will this create a problem?

The short answer is no. The ADO will use addresses for deployed networks which are sub-netted from those allocated to the fixed network. The connectivity to deployed networks will be over long-haul networks (or bearers), so that the routing path will be from Australia, even if the deployed forces are in a different region.

If a deployed ADF network needs to connect to a network belonging to a coalition partner, which could have addresses from a completely different range, this is not a problem. It is likely that an exterior gateway routing protocol (e.g. BGP) will be used at the boundary. This will need to be configured appropriately, it will normally be necessary to avoid a situation where, for example, traffic from the deployed ADF unit to a destination in Australia is routed over the ally's networks to their home nation and thence to Australia. Similar issues apply today with routing in IPv4, and will be solved in the same way in IPv6, by careful and intelligent router configuration.

If a deployed ADF network uses IP transit services from a coalition partner, or commercial ISP, then it is most probable that tunnelling will be used. This separates the routing domain of the ADO from that of the service provider, so again no addressing problem should arise.

Address Space Conclusion

The three worked examples of the previous section suggest that a wide-range of network sizes could be realised using a 6 level hierarchy. The example analysis shows that the network size can vary quite dramatically between 34 bits (/30 address) and 46 bits (a /18 address).

Because we have specified a lifetime of 15 years and we assume that we are more likely (at this point in time) to have under-estimated the potential for growth over that period, it is recommended that the ADO apply for a minimum allocation of a /18 address.

Ipv6 Transition Governance

This section details the recommended ADO IPv6 transition governance structure that should be used to manage and coordinate all the ADO's IPv6 transition activities.

Introduction

This section is introduced by revisiting our important high-level network centric principles and by providing some background to the potential difficulties that have been experienced as the result of other organisations approaches to their information environment governance structures.

As a starting point, the ADO's DIE governance structure must support the ability of the various ADO organisations to work together to achieve network-enabled operations by following our two (previously espoused) crucially important principles (see section 3.1.3 for further explanation of the principles):

Principle 1 : Unit-Level LANs

End-systems (e.g. sensors, weapons, Allies etc) are connected to "the network" and not to each other. They are attached to unit-level LANs which are in turn connected via a router to either a radio-WAN or a terrestrial WAN.

Principle 2 : Routable WANs

Make Radio-WANs and terrestrial WANs routable.

In general other organisations have tended to cast their CIOs into one of two roles, or in some cases the job description is a mix of these two roles, i.e.:

- a) as the program manager for the implementation of various information infrastructure projects, with responsibility for their budget and schedule, or
- b) as the interoperability custodian across a diverse range of projects and programs, some information environment related and some not (i.e. end-systems and platforms).

Both roles introduce major challenges, especially if the role encompasses responsibilities as both a program manager in the information environment space and as the interoperability custodian across the whole defence environment, including for the end-systems and platforms.

If a CIO is cast with only program manager responsibilities (e.g. hypothetically, DWACN, JP2072 and others) it is likely that:

- a) The CIO will have the potential to be successful at achieving **Principle 2**, but because their responsibility does not extend to the end-systems, it will be difficult to achieve **Principle 1**. Whilst the result may be a highly capable information network (the plumbing, routers, servers, cable etc), it is highly likely that the desired capability of network centric operations will not be realised to the extent required by the ADO. In this context the fact that the network is IPv6 capable largely becomes irrelevant.

The OSD CIO in the US DoD has gone down this path and has attempted to solve the resultant problem by splitting its organisation into a CIO's office, who looks after standards compliance, and the Networks and Information Integration (NII) office who is the networking advocate. This has however resulted in the CIO becoming the advocate for the very programs that he's supposed to have oversight for. There is every possibility that this could create many conflicts of interests with the result being less than fully successful.

We recommend that the optimum situation will be formed by instituting a governance structure (for IPV6 transition) that focuses on the CIO being the "interoperability custodian" where:

- b) The CIO has measures in place to ensure that **Principle 2** is applied by the “information infrastructure⁵³” projects managers and **Principle 1** is applied by all other project managers. These other project managers will be delivering projects that connect to the DIE in some way, they will be the end-systems and will include the platform projects (e.g. JSF, AWD etc). As long as the existing projects/programs are suitably structured and of a manageable size, then it is recommended that their program structures be left intact. The important concept is to put measures in place that allow the interfaces between projects to become compatible and interoperable.

Ideally the sequence of events in the process that should be adopted to achieve the optimum result is: interoperability, followed by modularity and modularisation followed by standardisation.

Management and Organisational Structures

Figure 23 illustrates the stakeholder organisations from an ADO IPv6 transition perspective.

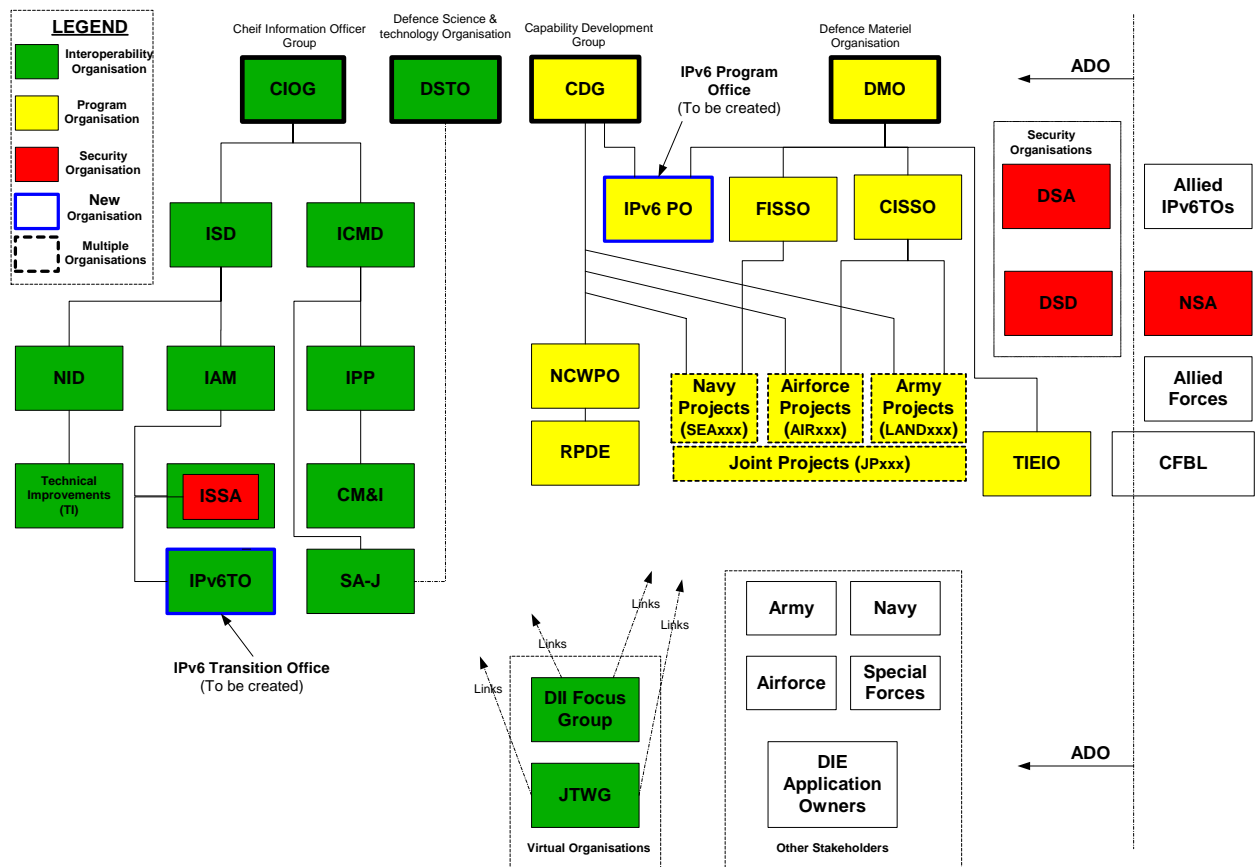


Figure 23 ADO Stakeholder Organisations From an IPv6 Perspective

The roles of the existing lead organisations and their existing subordinate organisations are described as follows.

Chief Information Officer Group (CIOG)

The CIOG is divided into the Information Systems Division (ISD) and the Information Capability Management Division (ICMD), for a complete CIOG organisational structure please see Annex F. Within these two divisions the following branches will have an IPv6 role as follows.

⁵³ e.g. DWACN, JP2072 etc.

Network Infrastructure Development (NID) Branch

As part of ISD, NID is an important organisation from an IPv6 perspective as responsibility for all the ADO's software applications have recently been centralised within the branch.

Technical Improvements (TI)

As part of NID, TI is expected to play a role in the development of an ADO pilot/test-bed⁵⁴ capable of implementing IPv6 for the purposes of evaluating the technology and the strategies for transition.

Information Architecture & Management (IAM) Branch

As part of ISD, IAM Branch is responsible for developing and maintaining the enterprise architecture and governance processes and tools that support the Defence Information Environment. Using its specialist staff and innovative support arrangements, IAM Branch assists Defence's Groups and Services in establishing and supporting their individual architecture offices and practices within the federated approach mandated by the Defence Architecture Framework.

Information Systems Security Assurance Branch

As part of IAM, ISSA will have responsibilities for the accreditation of applications and systems to the ADO IPv6 standard.

Information Policy and Plans (IPP) Branch

As part of ICMD, IPP branch executes the CIO's principal responsibilities as Coordinating Capability Manager of the Defence Information Environment (DIE). The Branch is responsible for the management and coordination of the DIE capability on a short-to-mid term basis (typically 0-5 years). The Branch is responsible for the short-to-mid term prioritisation of the information capability investment program (including minors) and oversight of portfolio DIE expenditure.

Scientific Advisor - Joint (SA-J) Branch

As part of ICMD, The Scientific Advisor - Joint (SA-J), represents the Defence Science and Technology Organisation (DSTO). SA-J advises the Australian Defence Joint Warfare, Information, Intelligence and Strategic communities on science and technology (S&T) issues and trends relevant to the development of capability and conduct/support of operations.

Defence Science and Technology Organisation (DSTO)

DSTO provides the ADO with scientific advice and supports the CDG and DMO through providing specialist scientific reports and conducting risk-analysis work and experiments in the support of these organisations.

Capability Development Group (CDG)

CDG is the ADO's capability manager and is responsible during the start-up phase of projects through to the completion of second-pass where the projects are handed over to the DMO. In Figure 23 we are showing this relationship for the Navy, Airforce, Army and Joint projects where there are links back to both CDG and DMO.

Network Centric Warfare Program Office (NCWPO)

The NCWPO has been established within the Integrated Capability Branch of CDG where it has authority to integrate projects into the force in being and the future force. The NCWPO is expected to be closely involved with ensuring that Principles 1 and 2 are followed.

Rapid Prototyping Development Environment (RPDE)

The Mission of the RPDE Program is "To enhance ADF war fighting capacity through accelerated capability change in the Network Centric Warfare (NCW) environment". The RPDE concept aims

⁵⁴ This was suggested by the Commonwealth at the IPv6 workshop. The hardware for this test-bed may already be in existence.

to create a collaborative, non-competitive environment where Defence and industry can seek opportunities where rapid enhancement to capability can be achieved, principally by incremental enhancement of existing capability.

Defence Materiel Organisation

Fleet Information Systems Support Organisation (FISSO)

The FISSO takes responsibility for Navy projects.

Command & Intelligence Systems Sustainment Office (CISSO)

The CISSO takes responsibility for Airforce and Army projects.

Tactical Information Environment Integration Office (TIEIO)

The TIEIO provides a support service to the DMO where it performs integration services for the ADO's tactical information environment including Tactical Digital Information Links (TADILs), e.g. Link-11, Link-16 and Link-22 etc.

Other ADO Stakeholders

DII Focus Group

This group currently does not exist. It should be formed as a virtual/matrix organisation of existing O6/EL2 level ADO members who will be responsible for leading the detailed planning phase, for making the required executive decisions in support of this IPv6 Transition Plan and tasking the JTWG (see 6.2.5.2). The DII Focus Group will have wider DII responsibilities than just IPv6. It should be noted that there is currently a Tactical Gateway Focus Group and it is recommended that this group is subsumed into the DII Focus Group⁵⁵.

Joint Technical Working Group (JTWG)

This group currently exists and should be expanded to receive IPv6 related tasking by the DII Focus Group. The JTWG will undertake further IPv6 related analysis and technical work to determine solutions and make more detailed proposals. It should be noted that at the time of writing this plan, the CIOG/IAM organisation will soon assume sponsorship and chair of the JTWG⁵⁶.

Security Organisations

The ADO's Defence Security Authority (DSA) is the ADO's internal security authority with oversight over security for the whole of the ADO.

The ADO's Defence Signal Directorate's (DSD) purpose is to support Australian Government decision-makers and the ADO with high-quality foreign signals intelligence products and services. DSD also directly contributes to the military effectiveness of the ADF, and provides a range of information security services to ensure that their sensitive electronic information systems are not susceptible to unauthorised access, compromise or disruption.

The ADO's IPv6 transition organisations (see section 6.3) will need to closely interact with both DSA and DSD on the security aspects of the transition from IPv4 to IPv6.

Others

The remaining stakeholders within the ADO fall into the category of users of the DIE and DII and the owners of applications that reside within the DIE. These users will be the subject of IPv6 communications (see section 6.3) and training programs.

⁵⁵ These recommendations are in accordance with Commonwealth comments to the draft IPv6TP.

⁵⁶ This was advised by Commonwealth comments to the draft IPv6TP.

Other Stakeholders

Stakeholders external to the ADO include:

- the Allied IPv6 organisational elements (e.g. US IPv6TO),
- the US National Security Authority who the ADO will need to interact with to achieve interoperability and
- other government organisations.

IPv6 Transitioning Organisations

To support the goal of providing a governance structure which:

- i) champions interoperability through policy measures and ensures that **Principles 1 and 2** are implemented and
- ii) avoids an organisational structure (from an IPv6 transition view) like the US DoD CIO and NII offices,

We recommend that two new organisations are created as indicated in Figure 23.

The first organisation, the IPv6 Transition Office (IPv6TO), will sit within the CIOG and will provide the CIO with the governance measures to ensure that the CIOG becomes the “interoperability custodian” for the transition of the ADO’s DIE and its end-systems from IPv4 to IPv6 i.e. to support **Principles 1 and 2** being implemented by CDG and DMO.

The second organisation, the IPv6 Program Office, will be functionally responsible to the CDG for projects during the start up phase and to the DMO for projects post second pass. It is enabled (by way of budget and schedule responsibility) to actually implement **Principles 1 and 2**. It is recognised that such an arrangement may be difficult to achieve and the resolution may be to create one program office within CDG and a second within DMO.

The roles and responsibilities of these two organisations are provided in more detail in the following sections.

IPv6 Transition Office (IPv6TO)

The prime function of the ADO’s IPv6TO is to operate as the “interoperability custodian” for IPv6 transitioning activities (please see section 7 for the recommended office organisational structure and position descriptions).

The ADO IPv6 Transition Office will be established to carefully plan and manage, at the enterprise level, Defence’s transition to IPv6 and will document this planning in the ADO IPv6 Transition Plan. This plan will be developed through broad consultation with key stakeholders. The ADO IPv6 Transition Office will be responsible for co-ordinating transition planning, analysis, testing and implementation efforts across Defence, promoting knowledge sharing, ensuring needed infrastructure is provided, and implementing a systematic program of outreach within Defence. The office will ensure that critical enterprise transition issues are prioritised and addressed.

The Transition Office will be responsible for providing policy and technical guidance to services, groups and projects/IPTs, and for defining procedures for approval and testing of migration/transition implementations.

The Transition Office should be structured to provide direction and guidance in the following technical areas:

- **Security:** this is a critical area; procedures must be in place to ensure that policy is enforced. System accreditors must be engaged to ensure that IPv6 issues are understood. There should be close co-ordination with ADO defence security organisations.

- **Networks:** this will cover both WANs and LANs. The provision of IPv6 service by individual component networks within the DIE will need to be co-ordinated. Issues relating to the provision and management of 6 over 4 and 4 over 6 tunnels and/or dual stack operation will need to be resolved. It will be important to reach agreement on where the responsibility for inter-working lies at network boundaries. This will probably need to be determined on a case-by-case basis.
- **Address allocation:** the Transition Office will be responsible for the management of IPv6 address allocation and the naming and addressing policy. It will also be responsible for the establishment of a root IPv6 Domain Name service (DNS). In performing these tasks the Transition Office will liaise with the Network Architecture Office.
- **Applications:** initially this area will be concerned with the provision of application layer gateways between IPv4 systems and IPv6 systems. Subsequently it will be important to address migration of applications to IPv6. This will apply to common services (e.g. e-mail) as well as specific applications. Over the long term, this may become the major effort in the transition process.
- **Allied interoperability:** the Transition Office should provide a central focus for discussions with Allies (principally the US DOD) on the co-ordination of IPv6 migration where necessary for interoperability.
- **Scheduling:** the overall schedule for IPv6 migration will be maintained by the transition Office, which will need to co-ordinate the schedules of DIE component networks and information systems.
- **Standards:** IPv6 is currently a general term used to describe a wide range of technical standards. The IPv6TO will be responsible for defining the IPv6 standards baseline for the ADO.
- **Testing:** it will be necessary to set technical specifications and standards to ensure inter-working of DIE components as they migrate to IPv6. The Transition Office will produce high-level test plans and have oversight of the testing process, this will include interoperability testing and IPv6 certification. The IPv6TO will also set criteria for the assessment of performance and inter-working.
- **Test-bed:** it is recommended that the ADO commence migration with a pilot implementation, in order to gain understanding and confidence before going forward to migration on operationally critical systems. The ADO IPv6 Transition Office should have close oversight of this pilot project. This pilot test-bed could either be newly constructed specifically for the purpose or hosted on one of the existing ADO test-beds⁵⁷.

IPv6 Program Office (IPv6PO)

The IPv6PO's prime function is to act as the Program Manager for the implementation of IP and the implementation of the transition of IPv4 to IPv6, ensuring **Principle 2** is implemented. This program level responsibility will extend to end-systems and platforms (outside the scope of the DIE), where the role is to ensure that **Principle 1** is implemented. As such the IPv6PO will have allocated budget to carry out its duties and will have schedule responsibility.

As the IPv6PO is expected to have minimal staff, individual projects/IPTs (Navy, Army, Airforce and Joint) will be required to contribute staffing resources to help the development of the Transition Plans, particularly in the area of cost and schedule estimates. The IPv6PO will work with the IPv6TO who will provide guidance and consultancy to assist the projects/ IPTs and to ensure reasonable consistency in the estimating process. The IPv6PO will have the following responsibilities:

- Program level responsibility (budget and schedule) for the implementation of IPv6 across all the ADO's projects.

⁵⁷ This issue was discussed at the IPv6 workshop. There are many test-beds within the ADO, 28 in ISD alone, there are J-series message test beds in the TIEIO and another 6 test-beds in the RPDE. There is also the ADO's involvement in the CFBL test-environment.

- High-level participant on the IPv6 Detailed Planning Phase in collaboration with the IPv6TO.
- Development, management and maintenance of the ADO's IPv6 Implementation Project Plan including cost and schedule.
- Responsibility for complying with the IPv6 governance measures and technical standards as set by the IPv6TO.
- Take direction for IPv6 related implementation tasks from various CDG and DMO managers.
- Overall management responsibility for the IPv6PO at an organisational and program/technical level. Will monitor progress against schedule and budget.
- IPv6 implementation interface with all CDG and DMO projects. This duty will require the incumbent to liaise with all impacted projects and programs (DWACN, JP2072, SEA1442 etc) to construct and maintain an overall ADO IPv6 schedule with inter-program/project dependencies. This schedule will also extend to Allied and other government programs.

Relationships with other IPv6 Transitioning Bodies

The ADO IPv6TO should establish and maintain close links with transition management organisations in Allied national defence departments. The prime link should be to the US DOD and DISA. The CCEB can facilitate this linkage and links to similar bodies in UK and Canada. It is expected that the US will wish to deal with Allies in multilateral bodies, rather than through many bilateral arrangements. The IPv6TO should also establish and maintain links with other Australian organisations (including industry) to achieve a whole of Government approach to the transition of IPv6.

Hierarchy of Documents

This section proposes a hierarchy of documents to be used by the ADO to manage and coordinate the IPv6 transition activities. This IPv6TP is the top-level parent document. The IPv6TO will maintain this document with changes and additions as required. Other IPv6 transition planning documents (including project budgets and schedules) will be subservient to this plan as depicted in Figure 24.

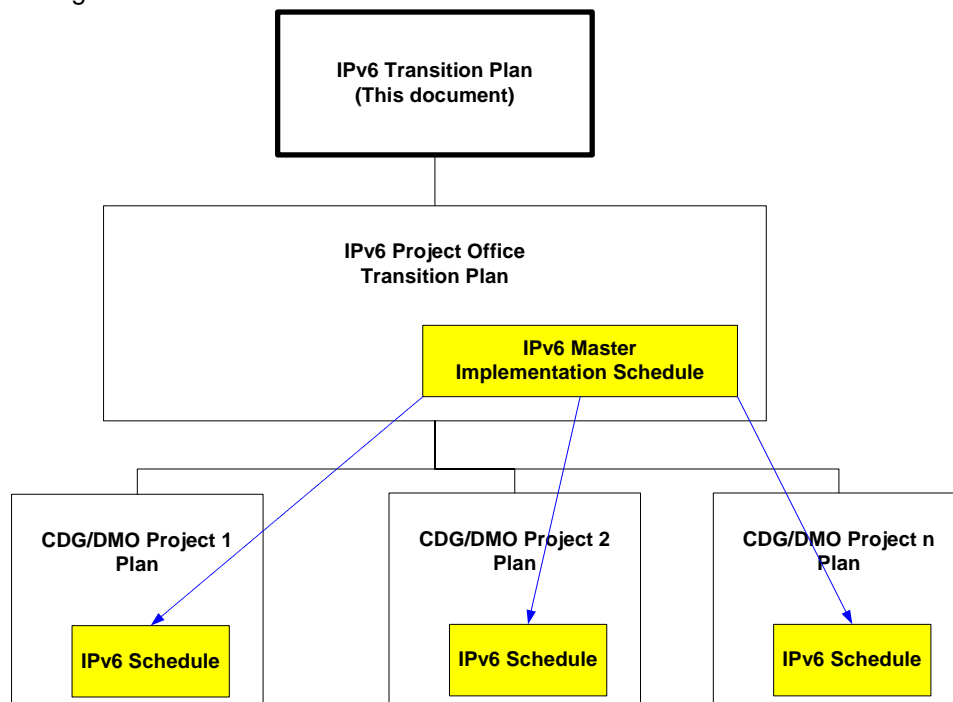


Figure 24 IPv6 Document Hierarchy

Lower Level IPv6 Transition Plans

The IPv6PO is assigned to coordinate the management of the implementation of IPv6 within the individual Defence Services, groups and projects. The IPv6PO will work closely with the IPv6TO and the projects being run by CDG and the DMO to develop an IPv6PO Transition Plan. A key function of the IPv6PO Transition Plan is to capture an integrated (whole of ADO) budget and schedule for the implementation of IPv6. The development of the top level integrated schedule (and budget) and the derived lower level (per project) schedules will require significant coordination and iteration between the IPv6PO and the individual projects. Also, because the transition strategy is leveraging off normal technology refresh cycles, these plans will require continual maintenance into the future.

These lower-level plans will be consistent with the overarching ADO IPv6 Transition Plan but will be focused on the planned transition within the Service/Group/Project, and will identify Service/Group/Project-specific issues and how they will be addressed. Critical dependencies and disconnects will be identified and worked through the DII Focus Group and the JTWG as appropriate. Individual Service/Group/Project plans would be endorsed by the ADO IPv6 Transition Office and approved within the individual Service/Group/Project.

These lower-level Transition Plans will include schedule and cost information. As they are produced it will become possible to refine the overarching ADO IPv6 Transition Plan, and the co-ordination process will lead to revision of these plans as necessary.

Ado ipv6 workforce requirements

To effect the ADO's transition to IPv6 over the period from now until 2013 will require effort to be applied in three major areas:

- i) Level of effort undertaken by staff within the IPv6TO,
- ii) Level of effort undertaken by staff within the IPv6PO and
- iii) Results/milestone based effort undertaken by other ADO staff (or contractors) to transition the DII's applications (software) and hardware.

The following sections propose a suitable workforce to cover the above areas of effort.

ADO IPv6 Workforce

The IPv6TO and IPv6PO will need to perform a number of functions that can be allocated to one or more of the respective office's staff members.

IPv6TO Functions

The functions to be performed by the IPv6TO include:

- i) Management and update of the ADO IPv6 Transition Policy [1].
- ii) Planning, management and implementation of IPv6 governance measures and processes.
- iii) Management of the transition of all DII applications and hardware from IPV4 to IPv6.
- iv) Management of the IPv6 Test Program, this includes management and oversight of an IPv6 Test-bed.
- v) Management of the IPv6 Security Program.
- vi) Management of the IPv6 Allied Interoperability Program.
- vii) Management of the IPv6 Communications Plan.
- viii) Definition and management of the ADO's IPv6 standard.
- ix) Provision of IPv6 technical specialist services.

IPv6TO Organisational Structure

The above functions of the IPv6TO could be fulfilled by an organisation with between three and four full time positions. The IPv6TO Lead may be a part-time (50%) position.




Figure 25 Suggested IPv6TO Organisational Structure

Lead Position

Position Duties

The IPv6TO Lead will perform the following duties:

- i) Take direction for IPv6 related tasks from various CIOG managers and the DII Focus Group.
- ii) Overall management responsibility for the IPv6TO at an organisational and program/technical level. Will monitor progress against schedule and budget.
- iii) Planning and management of the IPv6 governance measures and processes. This will include involvement in First and Second Pass project review processes from an IPv6 perspective.
- iv) Management of the IPv6 Communications Program. This duty will involve planning and performing IPv6 related education and information dissemination initiatives throughout the ADO, Allied organisations and other government organisations. The aim of the Communications Program is to ensure that the level of awareness within the ADO is sufficiently high across all the impacted ADO organisations to ensure an astute and timely transition from IPv4 to IPv6.
- v) IPv6 Coordination. This duty will require the incumbent to liaise with all impacted projects and programs (DWACN, JP2072, SEA1442 etc) to construct and maintain an overall ADO IPv6 schedule with inter-program/project dependencies. This schedule will also extend to Allied and other government programs.
- vi) Management of the IPv6 Allied Interoperability Program. This duty will involve planning and performing the various initiatives required to support Allied interoperability from an IPv6 perspective. The incumbent will be responsible for liaising with Allied IPv6 transition offices and ensuring that Allied related information is passed onto ADO projects as well as putting the case for ADO IPv6 requirements to Allies.

Owning Organisation

Each IPv6TO position is part of the CIOG organisation.

Competencies and Qualifications

The incumbent will need to be a competent manager of technology in a defence environment and is expected to hold a minimum of a diploma or degree qualification in engineering and preferably with a specialty in communications and network engineering. They will be capable of understanding the technical issues of transitioning the DIE to IPv6 and directing the technical effort within the IPv6TO.

Experience Required

The incumbent will need to have several years experience managing related projects within the DIE. A broad range of experience will be required in area of the terrestrial (DWACN) and deployed/tactical networks.

Tenure

This position will be required for the life-time of the IPv6 transition.

Test Manager Position

Position Duties

The IPv6TO Test Manager will perform the following duties:

- i) Take direction for IPv6 related testing and related tasks from the IPv6TO Lead and from various CIOG managers and the DII Focus Group in consultation with the IPv6TO Lead.
- vii) Overall management responsibility for the ADO's IPv6 test program at the program and technical level. The incumbent will be responsible for ensuring that the required

“IPv6 test-bed” assets are in place within the ADO. It is expected⁵⁸ that the Combined Forces Battle Lab Network (CFBLNet, see Figure 17) will have a major role in the IPv6 test program.

- ii) Overall responsibility for the program to assess all the DII's applications and hardware for transition to IPv6. This includes managing and undertaking all the recommended assessment activities listed in section 7.2.

Owning Organisation

Each IPv6TO position is part of the CIOG organisation.

Competencies and Qualifications

The incumbent will need to be a competent manager of technology in a defence environment and is expected to hold a minimum of a diploma or degree qualification in engineering and preferably with a specialty in communications and network engineering. They will be very capable of understanding the technical issues of transitioning software applications and hardware within the DIE to IPv6.

Experience Required

The incumbent will have experience in the acquisition, test and acceptance of complicated hardware and software systems in the areas of communications and networking. It is preferable that they also have experience in the development and implementation of integrated hardware and software systems. It is also preferred that this experience extends across both the terrestrial (DWACN) and deployed/tactical environments.

Tenure

This position will be required for the life-time of the IPv6 transition.

Security Manager

Position Duties

The IPv6TO Security Manager will perform the following duties:

- i) Take direction for IPv6 related security tasks from the IPv6TO Lead and from various CIOG managers and the DII Focus Group in consultation with the IPv6TO Lead.
- ii) This position will be responsible for coordinating and liaising with the ADO's security organisation including the Defence Security Authority (DSA), the Defence Signals Directorate (DSD) and the Information Systems Security Assurance (ISSA) branch of the CIOG.
- iii) The position will also be responsible for liaising with other security authorities and administrations, most importantly the US and UK authorities.

Owning Organisation

Each IPv6TO position is part of the CIOG organisation.

Competencies and Qualifications

The incumbent will need to be a competent manager of technology in a defence environment and is expected to hold a minimum of a diploma or degree qualification in engineering and preferably with a specialty in communications and network engineering. They will be very capable of understanding the security issues with the transitioning of software applications and hardware within the DIE to IPv6.

Experience Required

The incumbent will have experience in assessing software and hardware systems from a security perspective in a defence environment. They will need to have prior experience working with

⁵⁸ As advised by the Commonwealth during the IPv6 Working Group meeting on 29 June 2005.

similar issues with the ADO's security organisations and it is preferable that they have experience working with at least one other external security organisation e.g. the USA's NSA. It is also preferred that this experience extends across both the terrestrial (DWACN) and deployed/tactical environments.

Tenure

This position will be required for the life-time of the IPv6 transition.

Technology Specialist

Position Duties

The IPv6TO Technology Specialist will perform the following duties:

- i) Take direction for solving IPv6 technical issues from the IPv6TO Lead and from various CIOG managers and the DII Focus Group in consultation with the IPv6TO Lead.
- ii) Provide IP specialist technical support to the IPv6TO and to the ADO as a whole.
- iii) IPv6 is currently a general term used to describe a wide range of technical standards. The incumbent will be responsible for defining the IPv6 standards baseline for the ADO.
- iv) Be actively involved in designing and performing IP related tests (general areas and security related) and assessment activities on the IPv6 Test Bed and the DII.
- v) Liase with software engineers to evaluate application code for compliance with IPv6 and assessment of the level of effort required to move an IPV4 application to IPv6. Perform the same function with the relevant hardware engineers for the transition of hardware to IPv6.

Owning Organisation

Each IPv6TO position is part of the CIOG organisation.

Competencies and Qualifications

The incumbent will be the prime technical point of contact for IPv6 within the ADO and as such they must in the first instance have a very good overall technical competency in the areas of communications and networking technology. They will have general knowledge of IPv6 and will over a short period of time become the ADO's subject matter expert in IPv6.

They are expected to hold a minimum of a degree qualification in engineering with a specialty in communications and network engineering.

Experience Required

The incumbent will have experience with the design and implementation of IP systems within the DII.

Tenure

This position will be required for the life-time of the IPv6 transition.

IPv6PO Functions

The functions to be performed by the IPv6PO include:

- i) Program level responsibility (budget and schedule) for the implementation of IPv6 across all the ADO's projects, within the realms of the Capability Development Group and the Defence Materiel Organisation.
- ii) High level participant on the IPv6 Detailed Planning Phase in collaboration with the IPv6TO.
- iii) Development, management and maintenance of the ADO's IPv6 Implementation Project Plan including cost and schedule.

- iv) Responsibility for complying with the IPv6 governance measures and technical standards as set by the IPv6TO.

IPv6PO Organisational Structure

The IPv6PO is conceived as an Integrated Product Team (IPT) with lines of reporting back through to CDG and DMO as required by the stage of the project (with an IPv6 requirement) being managed. The IPT would be staffed by members of both CDG and DMO and is estimated to be equivalent to one full-time position.

Although only consisting of nominally one new full-time position, the IPv6PO IPT will be supported by individual projects who will allocate resources to support the responsibilities of the IPv6PO (see IPv6 Project Managers below).

IPv6 Program Manager

Position Duties

The IPv6PO Program Manager will perform the following duties:

- i) Program level responsibility (budget and schedule) for the implementation of IPv6 across all the ADO's projects.
- ii) High level participant on the IPv6 Detailed Planning Phase in collaboration with the IPv6TO.
- iii) Development, management and maintenance of the ADO's IPv6 Implementation Project Plan including cost and schedule.
- iv) Responsibility for complying with the IPv6 governance measures and technical standards as set by the IPv6TO.
- v) Take direction for IPv6 related implementation tasks from various CDG and DMO managers.
- vi) Overall management responsibility for the IPv6PO at an organisational and program/technical level. Will monitor progress against schedule and budget.
- vii) IPv6 implementation interface with all CDG and DMO projects. This duty will require the incumbent to liaise with all impacted projects and programs (DWACN, JP2072, SEA1442 etc) to construct and maintain an overall ADO IPv6 schedule with inter-program/project dependencies. This schedule will also extend to Allied and other government programs.

Owning Organisation

Each IPv6PO position nominally reports back through to either the CDG or the DMO depending upon the stage of the subject project. As stated above, because of the difficulties with creating such a dual reporting structure it may be necessary to have two separate IPv6POs.

Competencies and Qualifications

The incumbent will need to be a competent manager of technology in a defence environment and is expected to hold a minimum of a diploma or degree qualification in engineering and preferably with a specialty in communications and network engineering. They will be capable of understanding the technical issues of transitioning the DIE to IPv6 and directing the technical effort within the IPv6PO.

Experience Required

The incumbent will need to have several years experience managing related projects within the DIE. A broad range of experience will be required in area of the terrestrial (DWACN) and deployed/tactical networks.

Tenure

This position will be required for the life-time of the IPv6 transition.

IPv6 Project Manager

There will be many IPv6 Project Managers spread across the range of Army, Navy, Airforce and Joint projects. The level of effort required to support each position within a project will vary between a small part-time role to a larger part-time role depending upon the scale of the IP implementation.

Each IPv6 Project Manager will be responsible to the IPv6 Program Manager and will have responsibility for either:

- the implementation of IP within their project, if they are DIE related, or
- the implementation of Principle 1 for end-system/platform projects.

Workforce to Transition Applications and Hardware

The depth of DIE analysis conducted in support of this IPv6TP has been insufficient to allow an accurate estimate of the total effort in man-years to transition the DIE's thousands of applications and hundreds of thousands of hardware items. What is provided here is a sequence of steps that needs to be followed during the detailed planning phase to formulate a quantifiable measure of the effort required. Management of this work will be the responsibility of the IPv6TO.

DII Applications

Each DII application needs to be assessed using the following steps:

- i) The first step is to define the set of applications that will need to connect to the IP network either now or into the future.
- ii) Applications are then categorised as either COTS or in-house/specialist developed. For those that are COTS, the vendor should be queried as to the IPv4/IPv6 roadmap. If the application is scheduled for IPv6 transition as part of the normal product development cycle then the additional level of effort (over that expended for any other product upgrade) is limited to that required to meet the conformance standards which verify that the application meets the ADOs IPv6 standard.
- iii) For COTS applications where IPv6 is not on the applications developmental roadmap or the date for delivery of IPv6 is too far in the future, the vendor should be requested to provide a price and schedule for inclusion of IPv6 specifically for the ADO. If the cost and schedule is acceptable to the ADO then the upgrade would proceed with the normal conformance testing process being applied.
- iv) For in-house or specialist applications where the ADO has ownership of the source-code and design documentation for the application, it should be possible for qualified software engineers to inspect the quality of the software design and implementation for transition to IPv6. This process will determine the level of effort required to perform that software upgrade including documentation and testing. The outcome of this process will determine the cost effectiveness of upgrading the application. If it is cost effective to upgrade the application then the upgrade should be undertaken and the software put through the same acceptance into service processes as any other IPv6 enabled application. If it turns out to not be cost effective then there are two options, the first is to continue to use the application (and therefore continue to provide IPv4 support) or to seek an alternative application that is IPv6 capable.
- v) Should the above steps not lead to an acceptable solution, the alternatives include maintaining IPv4 support for the application or potentially seeking an alternative application that is IP agnostic i.e. a web-based application where the IP requirement falls to the browser application.

DII Hardware.

The DII hardware will have many of the same issues as software applications, except that some hardware items (usually peripheral devices) will possess an embedded IP stack where the stack cannot be upgraded via software (e.g. printers). The steps and solutions are the same for software except that it is most unlikely that it will ever be cost effective to upgrade lower cost peripheral hardware devices. In these situations where the hardware item cannot be replaced, the

only solution is to continue using the device and provide IPv4 support. This then becomes an obsolescence issue.

Risk management

Risk Log

The completed risk log is included in Annex C. The remainder of this section summarises the results of the risk log.

Risk Mitigation Strategies

Risk mitigation strategies are included in the “Treatment Strategies” column of the risk log in Annex C.

Risk Summary

With cognisance of the ADO IPv6 Transition context (see 3.1), the risk analysis process generated a risk log containing twenty six (26) risks, the log includes contributions from stakeholders who attended the IPv6 Workshop and others generated by the IPv6 Panel. The risks were rated for “Likelihood” and “Consequence” using the process from the ADO’s Project Risk Management Manual, a summary of the outcome of the rating process for these risks is provided in the matrix in Figure 26.

Likelihood Rating	Almost Certain					
	Likely			1	1	Extreme
	Possible		2	5	10	
	Unlikely			Medium		2
	Rare					1
		Insignificant	Minor	Moderate	Major	Severe
		Consequence Rating				

Figure 26 Risk Ratings Summary Matrix

As Figure 26 indicates, there were no “Extreme” level risks identified, fourteen (14) “High” level risks, eight (8) “Medium” level risks and four (4) “Low” level risks. The most common “Likelihood” rating was in the “Possible” category where 19 of the risks were classed. The most common “Consequence” rating was in the “Major” category where 11 of the risks were classed, although there was almost a 50/50 split between risks classed from “Insignificant” to “Moderate” and those in the “Major” to “Severe” category.

Each risk was also classified for the “Sources of Risk”⁵⁹ using the standard list of sources from the ADO’s Project Risk Management Manual. As this IPv6TP is the very first step of a transition activity that is currently scheduled to run over the next eight years until 2013, the results for the most common sources of risk highlight the critical importance of the governance structure and

⁵⁹ See Annex C for the complete list of Risk Sources.

controls employed by the ADO to effect the transition from IPv4 to IPv6. The most common sources of risk (in order of frequency) were:

- “Management activities and controls” followed by,
- “Technology and technical issues” followed by,
- “ADO project offices”.

Whilst much of the risk is sourced internally within the ADO, the reliance on COTS to effect the implementation of the transition and the need to interface with external bodies that lie outside the governance structure (e.g. Allied IPV6 transitions) means that there is also a large body of risk that lies outside the direct control of the ADO. The reliance on COTS is reflected by the next two most common sources of risk being in the classified areas of “Defence contractors” and “Maturity of technology required”.

Therefore the most sensible (and likely best value for money) mitigation/treatment strategies will concentrate on maintaining flexible governance structures, plans and technical architectures to allow the ADO implementation to cope with COTS IPv6 and Allied transitions that fail to meet the expected (and planned) timelines and budgets. A large part of this flexibility will be achieved by the wide-spread (across the ADO) adoption of Principles 1 and 2 (see Section 6.1).

Therefore a large part of the risk mitigation activity will be effected by the ADO (IPv6TO) periodically and continually revisiting this plan and maintaining a flexible stance to ensure that it can compensate for the effects of any realised risks during the transition from IPv4 to IPv6.

Dependencies and Key Assumptions

Key Assumptions

The key assumptions used to compile this document are as follows:

- **Ubiquitous IP** : Although it appears that it is not yet specific ADO policy to include in the architectural baseline requirements specifying that all Networks / Data-links / Bearers within the DIE become routable by implementing IP, it is a key assumption that the requirements for NCW and the general push toward maximising the usage of COTS will mean that system designers will consider IP as a candidate technology wherever possible. Further to this it is also assumed that IP will be the first “Layer 3” technology considered by system designers and if it is not chosen for DIE system past 2013, it will be because of other reasons e.g. cost, interoperability or performance.
- **IPv6 Program Synchronisation** : Although the ADO will liaise and coordinate with Allies and their programs to transition from IPv4 to IPv6, the coupling between these programs will be loose and not necessarily synchronised. This implies that the ADO must be prepared to inter-operate with Allies using both IPv4 and IPv6 for an extended period, probably well past 2013 and potentially up until any IPv4 flag day⁶⁰, should one be pronounced.
- **Security** : The security mechanisms in place today within the DIE and Allied networks use a mix of physical separation and encryption at the Data Link layer (2) or Network layer (3). Despite this there are techniques in place that allow a certain degree of interoperability between the ADF and its Allies. Ideally though, the most flexible, powerful and interoperable networks would be achieved if the DIE and Allied networks completely progressed to implementing the end-to-end network model with all security implemented at the Application layer (7), or “object-based”. As object-based security is yet to be mandated for the DIE and there is significant momentum in the DCP toward expanding the current security mechanisms (e.g. \$50 mil JP 2069⁶¹), it is assumed that there will be no fundamental change to the DIEs security architecture up until 2013.

⁶⁰ It is the Panel’s view that an IPv4 Flag Day is almost certain to never occur.

⁶¹ High Grade Cryptographic Equipment.

Conclusions

The ADO issued the policy “Transition To Internet Protocol Version 6” [1] in February 2005, this policy requires the DIE to transition to IPv6 by 2013 and importantly the policy states that IPv6 is an enabler for the ADO’s vision of NCW.

This IPv6TP has been developed by BSG in collaboration with a Panel (“the Panel”) of IPv6 subject matters experts from the IPv6 Forum, QinetiQ, the Naval Post Graduate School and the W2COG over the period from May through to July 2005. A draft of this plan was discussed at a workshop on 29 June 2005 and was attended by members of the Panel and various Commonwealth members. This Plan is considered to be a living document that will require revision and maintenance in order to keep pace with the rapid changes in networking technology. The scope for this IPv6TP includes the whole of the ADO’s DII and DIE.

Although the ADO’s IPv6 Policy was in place at the start of this task, the “context” for Internet Protocol (IP) (and the transition from IPv4 to IPv6 specifically) within the DIE was not apparent. Therefore the Panel’s first response was to use a top-down system engineering methodology and develop the “context”, this is the focus for Section 3.

The methodology in Section 3 analysed transitioning from artisan-based to industrial based information systems and then developed a definition for the GIG. The key observation from this analysis is that “modularisation” is the key to achieving interoperability (Note: Standards are also important but not key). To achieve modularisation (and then “net centricity”) the following crucial overall design principles were generated:

Principle 1 : Unit-Level LANs

End-systems (e.g. sensors, weapons, Allies etc) are connected to “the network” and not to each other. They are attached to unit-level LANs which are in turn connected via a router to either a radio-WAN or a terrestrial WAN.

Principle 2 : Routable WANs

Make Radio-WANs and terrestrial WANs routable.

The analysis also produced derived design requirements for the end-systems (these connect to the DIE) and classified end-systems that comply with these requirements as “Good Network Citizens”. The analysis also defined performance requirements for radio-WANs and proposed potential candidates for COTS re-use, i.e. IEEE 802.x standards are recommended as prime-candidates for consideration, even though some of the standards (e.g. WiMAX) will require modification to suit military systems. Two methods of dealing with legacy technology (during the IPv6 transition) were also considered, Cocooning and Layer 7 gateways, of the two Layer 7 gateways are viewed as more useful.

The context setting analysis concluded with a definition of the boundary between the non-DIE and the DIE, this is important because the boundary often extends into the ADF’s platforms where many of the “legacy” issues will be encountered in the future. The use of the developed DIE IP context extended beyond just the technical (implementation) domain and was pivotal to the generation of the governance structure and workforce plan to support the transition from IPv4 to IPv6.

Section 3 also summarised the history and plans of the IPv6 activities being conducted by the UK MOD, NATO and the US DOD. Much of the detailed information concerning these plans was not available to BSG but should be available to the ADO through its Government-to-Government links. A list of known IPv6 documents is provided in 1.3.2 and the Panel is pleased to offer its assistance to the ADO with obtaining access to this material. It was concluded by the Panel that because IPv6 has yet to progress to a sufficient state (anywhere in the world) there are currently

no “off-the-shelf” strategies that could be applied to the DIE. In fact the implementation of the US DOD’s IP governance structure is viewed as containing a few lessons learnt and attributes that should be avoided by the ADO. As a result of this IPv6TP, the ADO is likely to be in advance of many organisations with regard to its IPv4 to IPv6 transition, and potentially better placed to meet its desired time-schedule if the governance mechanisms can be smoothly and successfully implemented.

The next input to the development of the IPv6 transition strategy was to analyse the current and future Defence Information Environment (DIE), see 3.3. The 2005 architecture and network magnitude was detailed with a specific emphasis on the DWACN, the DWACN is seen as a core element of the transition activity. This information was used as an input to the development of an IPv6 numbering plan in Section 5. The future DIE architecture was covered by specifying the DCP projects that will move the DIE from its current baseline to its future state.

Section 3 concluded by providing relevant challenges, opportunities and emerging technologies. The ADO can expect to find its major challenges in the areas of transitioning its non-routable networks and security. To ensure that the ADO can rise to these challenges, the IPv6TO is proposed to be staffed with positions (Technology Specialist & Security Manager) that specifically address these areas of challenge.

The recommended IPv6 transition strategy commenced in Section 4 by considering three options. A “big bang” strategy was deemed too risky and costly, an incremental approach with hard-milestones did not comply with the approach of leveraging off the natural technology refresh cycles and so this led to recommending an incremental transition with soft milestones. The recommended strategy is depicted in Figure 15, this shows seven overlapping (soft milestone) phases:

- Phase 1 Planning,
- Phase 2 Network Security,
- Phase 3 National Application Gateways & Allied Application Gateways,
- Phase 4 Overlay Networks,
- Phase 5 IPv6 Clouds,
- Phase 6 Cloud Expansion and
- Phase 7 End State.

Importantly this strategy allows for a progressive roll-out of IPv6 whilst recognising that some parts of the DIE may never transition and small enclaves of IPv4 and links to external IPv4 networks will be required past 2013. The Planning phase extends for the life-time of the transition and it will be the IPv6TO and IPv6PO who will be responsible for conducting this planning effort and maintaining the over-arching IPv6 documentation. Also, the Network Security, National Application Gateways and Allied Application Gateway phases will also span the entire transition period (although having staggered starts) due to their importance and need to iterate with changing conditions both within the DIE and those of influence external to it.

The strategy has also been designed to be cost-effective, to have no impact on defence operations and not to degrade interoperability with Allies, justification is provided in 4.3. To reduce the level of risk and ensure a successful transition Section 4.4 proposed a range of information assurance and test activities that will need to be conducted. These are designed to help ensure that ADO security is not prejudiced and experienced can be gained by the ADO with IPv6 before rolling the capability out into the DIE and operational environment. The recommended strategy section concludes with some specific advice for the DCP projects that are seen to be key to the IPv4 to IPv6 transition. Included in the list of key projects is JP2047, JP2008, JP2072, SEA1442 and JP2030.

At this early stage of the planning process it has not been possible to develop an IPv6 address space plan that can withstand the test of time i.e. from now until 2020 and beyond. However Section 5 provides a detailed step by step analysis method that can be used during the detailed

planning phase to construct a robust IPv6 address plan for the ADO. The method is illustrated using several examples, these are related to the ADOs current DIE architecture and consider future technologies that may be taken up by the ADO and consume addresses. These examples suggest that the ADO's IPv6 address range could be anywhere between 34 bits (/30 address) and 46 bits (/18 address). However the ADO should attempt to gain access to the largest contiguous block of addresses (e.g. /18) it can as the cost of using these addresses is likely to outweigh the costs of modifying the network in the future to suit a smaller (and or fragmented) address range.

Section 6 details a recommended governance structure for the ADO to transition the entire DIE and to ensure that the end-systems that attach to the DIE also conform to this IPv6TP. The "Unit-Level LANs" and "Routable WANs" principles (see above) are used again as the basis for the development of the governance structure. The CIOG organisation has recently undergone some significant changes and these are captured in the plan, the recent transition of the DMO to a prescribed agency may also have some impact on implementing these governance measures. Two new organizational offices are proposed to ensure that the governance regime is implemented in a astute and timely fashion and that the actual implementation of IPv6 is appropriately funded and scheduled.

The IPv6 Transition Office (IPv6TO) will be part of the CIOG, its prime responsibility will be as the "interoperability custodian" where it will complete the detailed planning of IPv6, promote information sharing across the ADO and ensure that the critical enterprise transition issues are prioritised and addressed. The IPv6TO will become the ADO's centre of excellence for IPv6 and will also offer technical guidance to the whole of the ADO. The IPv6TO will be staffed with up to four full-time positions.

The IPv6 Program Office (IPv6PO) has been proposed to act as the Program Manager for the implementation of IP in general and the transition of IPv4 to IPv6 across the whole DIE. Functionally the office must cover the scope of ADO projects from inception through to first pass (where they are under the control of the CDG) then on through second pass and into service (where they are under the control of the DMO). It is recognised that this may be a difficult proposition and if a single (one-person) office cannot be created, then the solution may be to have one office within the CDG and the other within the DMO. The IPv6PO will also require each project to allocate budget and schedule to the implementation of IPv6 as required. Although the office is small, its creation, function and lines of reporting are seen as crucial to a successful transition.

Section 7 details the organisational structure of the IPv6TO and IPv6PO. Each position within these offices is provided with a position description and description of the required competencies and experienced required to fulfil the role.

Although some quantification of the magnitude of the elements (hardware and software) of the current baseline DIE are provided in 3.3.1, it has not been possible within this IPv6TP to provide any detailed estimates for the level of effort required to transition software applications and hardware. Section 7.2 does however provide a detailed step by step procedure for assessing the hundreds of applications within the DIE with the aim of determining the effort/cost of making the IPv4 to IPv6 transition. It is also recognised that some applications may not be cost effective to transition and will be maintained as is in IPv4 enclaves within the DIE.

The conclusion to the process of developing this IPv6 transition strategy was to assess all its elements (including the proposed governance structure and workforce) for risk, see Section 8. A risk log capturing 26 risks was developed, each risk was assessed for likelihood and consequence and mitigation strategies were proposed. As the IPv6 transition will be heavily dependent upon COTS, there is a large degree of risk that will be beyond the direct control of the ADO. The responsibility for managing this risk will rest with the IPv6TO who will need to continually revisit this plan.

Recommendations

This IPv6TP is the first major step in an eight-year project to transition the entire DIE to IPv6 by 2013. As such the ADO will be required to complete many inter-linked activities and work with a variety of external organisations during the lifetime of the project.

The following is a list of recommendations for the immediate term:

- the ADO endorse this IPv6TP,
- the ADO endorse the governance and organisational components of this IPv6TP and commence resourcing the IPv6TO and IPv6PO,
- continue to engage the community of IPv6 subject matter experts to ensure that the progress with other organisations is tracked and lessons learnt are continually captured,
- sponsorship of combined Defence/Industry IPv6 forums to expand Defence's engagement with industry and whole of government,
- commence the Detailed Planning Phase including an initial IPv6 threat assessment and
- review the ADO's DWACN as-is and future architectures descriptions for the impact of this IPv6TP.

The following is a list of recommendations for the medium term:

- maintain and update this IPv6TP and its associated policies,
- undertake a detailed review each of the "Key Projects For Transition" (see 4.5),
- conduct a more detailed study and workshop in support of extending the work in this IPv6TP and developing a future looking IPv6 address plan and
- undertake specialist IPv6 and Network Centric focussed (See principles 1 and 2) training to raise the level of expertise within the ADO.

-

ANNEX A Interoperability Options

⁶²IPv6 and IPv4 will coexist for many years. A wide range of techniques has therefore been defined that make the coexistence possible and provide a path toward transition. These techniques fall into three main categories:

- Dual-stack,
- Tunnelling and
- Translation.

Dual stack techniques allows IPv4 and IPv6 to coexist in the same devices and networks. Tunnelling techniques allow the transport of IPv6 traffic over the existing IPv4 infrastructure. Translation techniques allow IPv6-only nodes to communicate with IPv4-only nodes.

Dual-stack Techniques

Using the dual-stack nodes throughout a network provides complete support for both IPv4 and IPv6 protocol versions.

In communication with an IPv6 node, such a node behaves like an IPv6-only-node, and in communications with an IPv4 node, it behaves like an IPv4-only node. Implementations probably have a configuration switch to enable or disable one of the stacks. Therefore dual stack nodes can have three modes of operation:

- IPv4 enabled and IPv6 disabled – Behaves like an IPv4 only node
- IPv4 disabled and IPv6 enabled – Behaves like an IPv6 only node
- IPv4 enabled and IPv6 enabled (IPv4/IPv6) – Node can use both protocols

An IPv4/IPv6 (both stacks enabled) node has at least one address for each protocol version. For IPv4 it will configure by using either static configuration or Dynamic Host Configuration Protocol (DHCP) and for IPv6 it will use either static configuration or auto-configuration.

Domain Name System (DNS) is used with both protocol versions to resolve names and IP addresses. An IPv4/IPv6 node needs a DNS resolver that is capable of resolving both types of DNS addresses records. In some cases, DNS returns only an IPv4 or an IPv6 address. If the host that is to be resolved is a dual-stack host, DNS might return both types of addresses. Generally, applications that are written to run on dual-stack nodes need a mechanism to determine whether it is communicating with an IPv6 peer or an IPv4 peer.

A dual-stack network is an infrastructure in which both IPv4 and IPv6 forwarding is enabled on all routers. The disadvantage of this technique is that a full network software upgrade is required to run the two separate protocol stacks. This means all tables (e.g. routing tables) are kept simultaneously, routing protocols being configured for both protocols. For network management, there are separate commands (e.g. Windows OS - ping.exe for IPv4 and ping6.exe for IPv6). Other problems include higher memory and power consumption.

Dual-Stack Advantages

- Easy and flexible to use.
- Hosts can communicate with IPv4 hosts using IPv4 or with IPv6 hosts using IPv6.
- When the IPv6 upgrade is complete the IPv4 stacks can simply be disabled or removed.

⁶² Most of this Annex has been sourced from [2] and the IPv6 Forum

Dual-Stack Disadvantages

- Two stacks require more CPU power and memory than one stack (not such a big issue).
- Requires two tables, one for each protocol, increased management effort.
- Requires two sets of commands, one for each protocol, increased management effort.
- A DNS resolver running on a dual-stack host must be capable of resolving both IPv4 and IPv6 address types.
- Applications on a dual-stack host must be capable of determining whether this host is communicating with an IPv4 or IPv6 peer.
- Should use a firewall to protect the IPv4 network and the IPv6 network.

Tunnelling

Tunnelling is used to carry IPv6 traffic by encapsulating it in IPv4 packets and tunnelling it over the IPv4 routing infrastructure.

Tunnelling

There are two types of tunnelling⁶³:

- Manually configured tunnels. IPv6 packets are encapsulated in IPv4 packets to be carried over IPv4 routing infrastructure. These are point-to-point tunnels that need to be configured manually.
- Automatically configured tunnels. IPv6 nodes can use different types of addresses (e.g. IPv4-compatible-IPv6 addresses, 6to4 or Intra-Site Automatic Tunnel Address Protocol (ISTAP)) to automatically tunnel packets over the IPv4 routing infrastructure. These special IPv6 uni-cast addresses carry an IPv4 address in some of the IPv6 address fields.

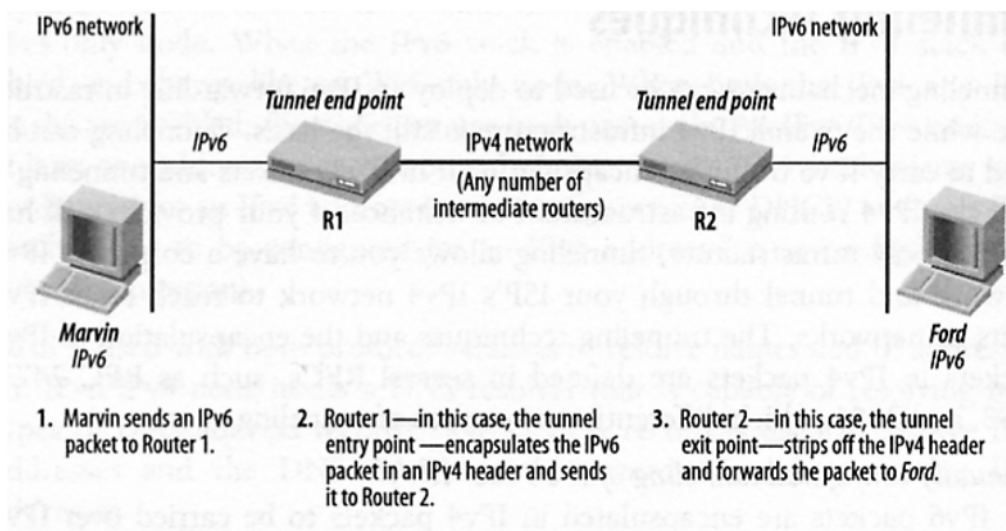


Figure 27 Tunnelling (6-over-4) Example

Tunnelling Advantages

- Flexibility, there is no specific upgrade order that needs to be followed.
- Single hosts or single sub-nets within a corporate network can be upgraded to IPv6.
- Continue to use IPv4 core network (Telstra & Singtel/Optus), core doesn't need to support IPv6.

⁶³ For more info see IETF RFC2893

Tunnelling Disadvantages

- Additional load placed on the router (a vendor design problem only).
- Tunnel entry and exit points need time and CPU power for encapsulating and de-encapsulating packets (a vendor design problem only).
- Single point of failure (can be overcome by better network design).
- More complex trouble-shooting as may develop “hop count”, MTU size or fragmentation issues.
- Less flexibility when using IPv4 compatible IPv6 address, as the limitations of the IPv4 address space remain in place.
- Potential for the number of tunnels to become very large and unmanageable.

Manually Configured Tunnels⁶⁴

A manually configured tunnel is an IPv4 or IPv6 tunnel configured between two end-points to carry IPv4 or IPv6 traffic. This allows for example two IPv6 networks to be connected even when the infrastructure between those two networks is not IPv6 capable, or later in the transition two IPv4 networks to be connected that are separated by an IPv6 network.

Advantages

- Simple to deploy inside a network
- Allows transport of IPv6 packets over an IPv4 network
- Available on most platforms
- Also supports IPv4 traffic over IPv6
- Permits end-to-end interoperability
- Permits end-to-end secure trust model
- IETF Standard and specified solution

Disadvantages

- Must be manually configured
- Due to management overhead does not easily scale to be used in end-hosts
- May not scale without automation for many users across routing fabric

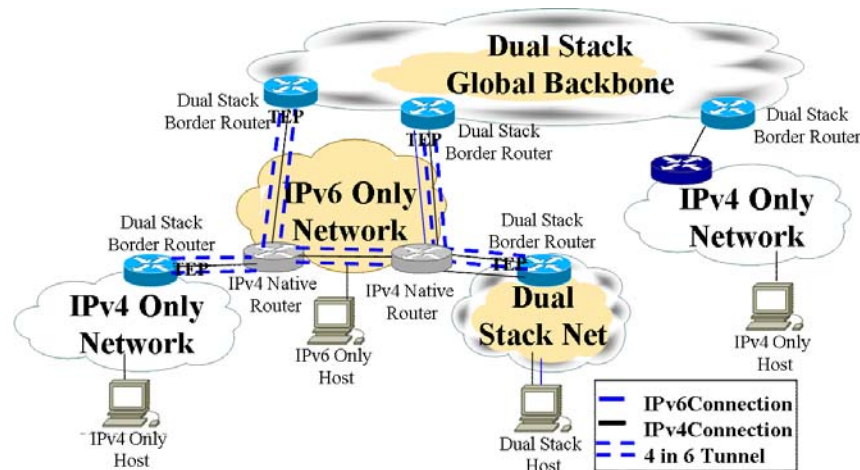


Figure 28 Manually Configured Tunnelling Example

⁶⁴ <http://ftp.rfc-editor.org/in-notes/internet-drafts/draft-ietf-v6ops-mech-v2-07.txt> This document obsoletes RFC 2893.

Automatically Configured Tunnels

The following automatic techniques, 6to4, ISATAP and Teredo are expected to be more applicable for the commercial domain where more flexibility is required, mostly because it is common for the other end of the tunnel to be beyond the control of the network administrator. For defence applications it is expected that manual tunnelling methods will be more appropriate because increased control is provided and defence will have “mostly” complete control over its infrastructure.

6to4⁶⁵

This is a mechanism that requires a single, globally unique IPv4 address. By embedding this 32 bit IPv4 address into a reserved IPv6 prefix, a router can create a globally unique /48 IPv6 prefix. The IPv6 packets are encapsulated in IPv4 packets without using explicit tunnels but automatic tunnelling mechanisms. Thus, making this low configuration overhead mechanism especially useful in IPv6 capable end-hosts. The usage of the special 6to4 address format, however, prevents the usage of an operator's own address space. Thus, 6to4 is impractical in roll-outs beyond single host configurations or very small networks.

Advantages

- Relatively easy to deploy.
- Supported on numerous platforms.
- Provides an address block for an AS without dealing with any registry.
- An existing standard (RFC 3056).
- Permits end-to-end interoperability.
- Permits end-to-end secure trust model.
- Public 6to4 relays exist today.

Disadvantages

- Operator's allocated IPv6 address space cannot be used.
- Impractical in network based roll-outs when entity has their own IPv6 prefix.




Figure 29 6to4 Example

⁶⁵ For more information see IETF RFC3056

ISATAP⁶⁶

ISATAP is a solution to provide IPv6 connectivity to sparsely located hosts in an IPv4 network. In this solution, the IPv6 capable hosts use automatic IPv4 tunnels to connect to an ISATAP router, which is connected to further IPv6 networks.

The ISATAP addresses are created by using a special ISATAP interface identifier derived from the host's IPv4 address. Once the ISATAP router is reached, standard IPv6 stateless

Autoconfiguration is used over the automatic tunnel to create the IPv6 address. This mechanism allows the usage of operator allocated IPv6 address space in the prefix. In addition, the hosts can be in private address space as long as there is not Network Address Translation (NAT) between the hosts and the ISATAP router. Intra-site communication can be done directly between hosts using the IPv4 address in the interface identifier.

Advantages

- Provides for easy incremental deployment of IPv6 to disparate nodes in a site.
- Supported on many platforms.
- Works in sites that use private addresses when NAT is not present.
- Permits end-to-end interoperability.
- Permits end-to-end secure trust model.
- IETF work in progress, but unknown if it will be standardised by any entity.
- Supported by some platforms
- ISATAP will self-deprecate (i.e. turn itself off) when IPv6 is dominant and in use without the network operators having to dismantle the ISATAP mechanisms.

Disadvantages

- Caution has to be used when deploying the ISATAP routers to make sure they are not used to hide a Denial of Service (DoS) attack.
- ISATAP does not provide for multi-cast support




Figure 30 ISATAP Example

⁶⁶ <http://ftp.rfc-editor.org/in-notes/internet-drafts/draft-ietf-ngtrans-isatap-24.txt>.

Teredo⁶⁷

Teredo provides IPv6 connectivity to hosts that are located behind NAT-devices in networks without IPv6 support. Teredo uses a special address format where the IPv6 prefix is created using special Teredo prefix, IPv4 address and a UDP port number. The IPv6 packets are encapsulated in UDP allowing NAT traversal. The IPv6 address is automatically configured to the Teredo host by a Teredo server in the Internet. Two Teredo hosts can also use direct tunnelling between themselves.

Advantages

- Easy to implement on a “one-off” basis.
- Provides a solution that works through NATs.
- Provides a solution for networks with no IPv6 support.
- IETF standardized solution in process

Disadvantages

- Uses a special IPv6 address format. Thus, operator’s own allocated address space cannot be used.
- Uses UDP to force hole in the client firewall.




Figure 31 Teredo Example

Tunnel Broker Overview⁶⁸

Tunnel Setup Protocol (TSP) is a tunnel broker based solution where the TSP client connects to a TSP broker that sends the client configuration information for setting up the tunnel between the TSP client and a tunnel server. TSP works both for a single host and a router. In addition, of providing IPv6 connectivity TSP supports authentication of the user and supports tunnelling of IPv4 over IPv6.

TSP is a good solution for connecting IPv6 networks as it supports IPv6 prefix delegation. In addition, TSP supports UDP encapsulation of the packets enabling NAT traversal of the tunnel. TSP can use operator’s own address range for the terminals.

⁶⁷ <ftp://ftp.rfc-editor.org/in-notes/internet-drafts/draft-huitema-v6ops-teredo-05.txt>.

⁶⁸ <ftp://ftp.rfc-editor.org/in-notes/rfc3053.txt> & <ftp://ftp.rfc-editor.org/in-notes/internet-drafts/draft-blanchet-v6ops-tunnelbroker-tsp-02.txt>

TSP has not been standardized in any standardization body, yet. However, there are activities on-going to bring TSP to the IETF.

TSP is an instance of the Tunnel Broker model (RFC 3053). The TSP allows authentication of the user at tunnel setup.

Advantages

- Smaller configuration overhead than manually configured tunnels.
- Works also in dynamic environments.
- Supports NAT traversal.
- Support tunnelling of IPv4 over IPv6.
- Supports DSTM (below).
- Referenced as method to review by U.S. DoD.
- Strong industry support for deployment

Disadvantages

- Large signalling overhead.
- Heavy solution.
- Not standardized yet, but in process.

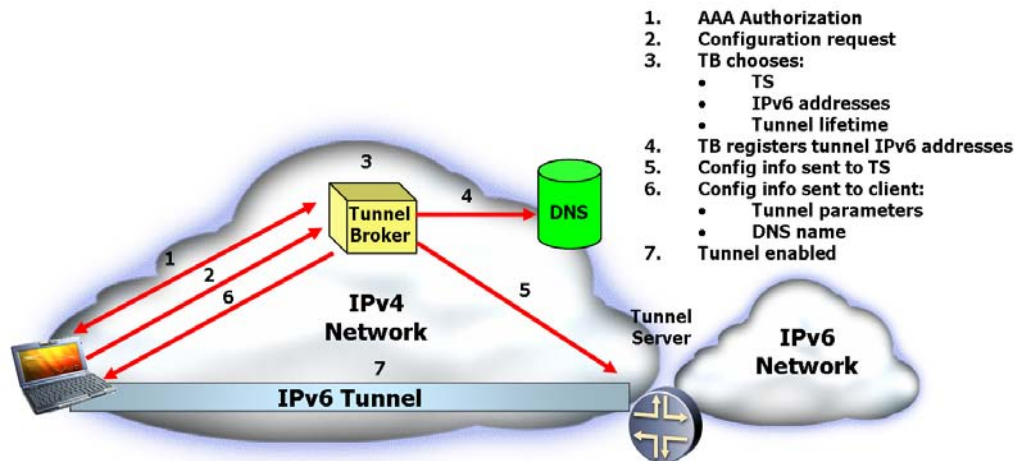


Figure 32 Tunnel Broker Example

Dual Stack Transition Mechanism (DSTM)⁶⁹

DSTM has a different assumption to transition than the other mechanisms. DSTM assumes an IPv6 dominant deployment where most of the hosts are IPv6 capable and the network is mostly IPv6 only. In DSTM, IPv4 is transported over IPv6 tunnel to an IPv4 network.

IPv6 deployment in some operational networks will use an IPv6-dominant network deployment strategy. What IPv6-dominant means is that the network will transition to IPv6 using only IPv6 routing to transfer both IPv4 and IPv6 packets.

Advantages

- Provides IPv4 connectivity in IPv6 networks without explicitly configured tunnels.

⁶⁹ <http://www.rfc-editor.org/cgi-bin/iddoctype.pl?letsgo=draft-bound-dstm-exp-03>

- Maintains end-to-end security for IPv6 connectivity and for IPv4, when enough IPv4 global address space is available.
- Has had industry implementation and some testing.
- Referenced as method to review by U.S. DoD.
- Strong Industry support for this method.

Disadvantages

- Not standardized, yet, but in process.




Figure 33 DSTM Example

Translation⁷⁰

Network Address Translation – Protocol Translation (NAT-PT) uses address translation. Basically NAT-PT translates IPv6 packets to IPv4 packets and visa versa. The NAT-PT device has to keep state information of the flows passing the device to perform the protocol translation. The mechanism relies on a DNS Application Level Gateway (ALG) to translate IPv6 address queries to IPv4 queries and to build up the state in the NAT-PT device. The usage of the DNS-ALG is seen problematic due to various reasons. Thus, the IETF is in the process of moving the NAT-PT standard to experimental RFC.

The NAT-PT solution allows IPv6 only nodes in an IPv6 network to communicate with IPv4 nodes without being directly connected to the IPv4 network. However, it does have the same shortcomings and restrictions than regular IPv4 NAT has. Thus, applications that do not work well with NATs do not work with NAT-PT either.

Translation Advantages

- Transparent to end nodes. Easily provide IPv4/IPv6 interoperability.
- Mechanism that allows the continued use of mission critical application or services that may be undesirable to have ported for use with IPv6.

Translation Disadvantages

- Single point of failure/bottleneck.
- Added administration.
- Has the same shortcomings of a traditional NAT.

⁷⁰ <ftp://ftp.rfc-editor.org/in-notes/internet-drafts/draft-ietf-v6ops-natpt-to-exprmntl-01.txt>

- DNS-ALG is seen problematic.
- Does not permit the end-to-end network model

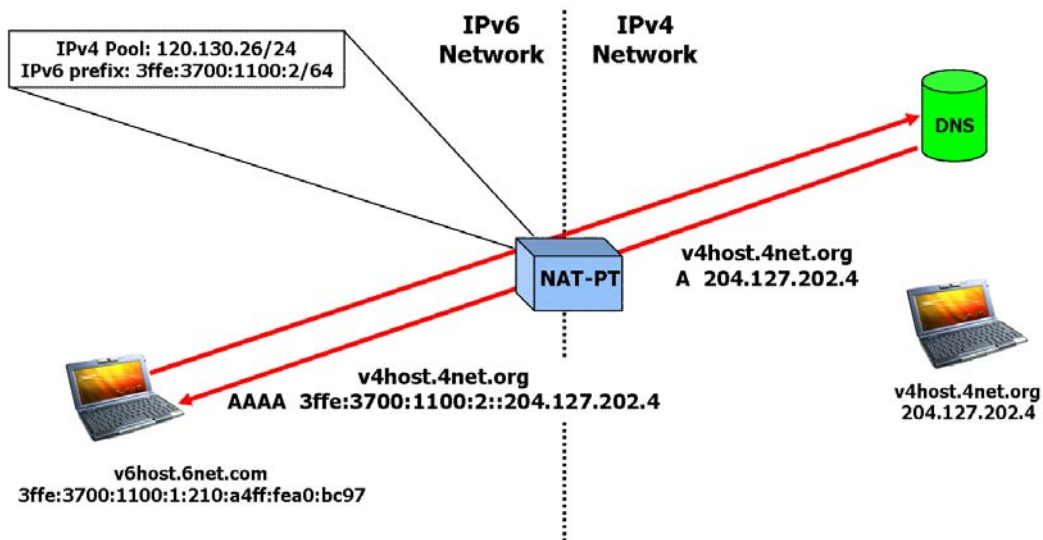


Figure 34 Translation Example

Annex B Phase 1 Detailed Planning

IPv6 Address Planning

Using the framework provided in Section 5 “IPv6 Address Space Requirements”, an address plan will be developed for the whole of the DIE.

Non-routable network Planning

Detailed planning for the DIE’s non-routable networks will consist of the following tasks:

- Determine ingress and egress points to the Non-Routable networks within the ADF DIE.
- Determine if interfaces can be specified with IP from those ingress and egress interfaces.
- Determine the data structures for those interfaces, and then a network proxy or gateway will have to be developed to support those input and output interfaces.
- Determine what semantics relative to the network are contained within those interfaces.
- If IPv4 is assumed in those interface semantics, then IPv4 should be used to input to this network.
- For an IPv6 Transition when a packet flows into or through a Non-Routable network the Transition should assume it should be presented an IPv4 packet, this implies potentially translating IPv6 to IPv4. This could have great cost and affect end-to-end interoperability for any IPv6 context and assumptions for security end-to-end.
- It would be best if possible to redefine the Non-Routable network interfaces to support IPv6 from the beginning if at all possible for the IPv6 Transition.

Interoperability Planning

Detailed planning to achieve the required interoperability will mostly be conducted by individual projects where they will need to perform a range of tasks including:

- Determine set of network applications⁷¹ that must be ported / invented.
- Determine the geography the network applications must span.
- Identify Network components that must support IPv6.
- Identify Network components that require IPv6 Transition Mechanisms.
- Identify Network components that can be initiated with IPv6 using IPv4 as scarce resources only.

Determine the packets required over the DIE within the scope of the IPv6TP:

- Packets over a local link, a site, an intranet, an Internet and over a mobile IPv6 network.
- Packets from IPv6 Network thru IPv4 Cloud to IPv6 Network.
- Packets from IPv4 Network thru IPv6 Cloud to IPv4 Network.

Determine the points of network communications for the IPv6TP Node Types:

- Clients, Servers, Routers, Switches, Printers, Gateways, Firewalls, Proxies, and any network device or applications platform.
- Management Nodes (e.g. Network, Security, Mobility, QoS).
- Any Node supporting Transition Mechanisms.
- Public Key Infrastructure Nodes for Security.

Determine the points of network communications for the IPv6TP Software Components:

⁷¹ In general it is expected that most applications within the DIE will need to migrate to IPv6, except for those completely stand-alone applications including those which do not connect to a LAN.

- Network Management and Utilities.
- Network Internet Infrastructure Applications.
- Network Systems Applications.
- Network End User Applications.
- Network High Availability Software.
- Network Security Software.

Costing

Once the above detailed planning is completed, each individual project should have a sufficient information and implementation level detail to complete the costing exercise for transitioning to IPv6.

ANNEX C Risk Log

#	Risk Author	Affected Components or Systems	Component or System Description	Description of Risk	Sources of Risk		Evaluation of Existing Controls	Likelihood of Risk	Consequences of the Risk	#	Likelihood Rating Value	#	Consequence Rating Value	Overall Rating (from DoD matrix)	Accept or Treat Risk? (with reasons why)	Treatment Strategies	
1	IPv6 Workshop (Grant Ranard)	DWACN	DWACN	The DIE ends-up without an overall end-to-end security architecture (there currently is no such end-to-end architecture)	6	7		Unknown/TBD	Possible: If the treatment strategies are not followed or fail.	Major: Because security is a key requirement for the DIE.	3	Possible	4	Major	High	Risk must be treated as this risk has such wide-ranging impact.	Provide budget for architectural work, undertake the work and manage the effort via governance and management structures.
2	IPv6 Workshop (Group)	HAIBE	HAIBE	There may be insufficient quantities of HAIBE IPv4 IPv6 (combinations) encryptors made available to the ADO.	1	5	7	Unknown/TBD	Possible: If the treatment strategies are not followed or fail.	Major: Because security is a key requirement for the DIE.	3	Possible	4	Major	High	Risk must be treated as this risk has such wide-ranging impact.	Raise the level of importance of the issue by appropriate use of government to government channels.
3	IPv6 Workshop (David Holmes)	All IPv6 affected	Accreditation	ISSA/DSD/DSSA delay/deny IPv6 accreditation.	7	19		Unknown/TBD	Possible: If the treatment strategies are not followed or fail.	Major: Because security is a key requirement for the DIE.	3	Possible	4	Major	High	Risk must be treated as this risk has such wide-ranging impact.	Ensure that these security organisations are suitably staffed.
4	IPv6 Workshop (Grant Ranard)	DIE	DIE	IPv6 policy implementation fails	7	19	3	Unknown/TBD	Likelihood will be a function of either failing to provide sufficient penalties (sticks) and rewards (carrots)	Severe: Because it is possible that the ADO incurs an obsolescence problem and or an interoperability (with allies) problem.	3	Possible	5	Severe	High	Risk must be treated as this risk has such wide-ranging impact.	Determine set of metrics to monitor during course of implementation. Use results to apply changes to policy to avoid failure.
5	IPv6 Workshop (Group)	Major systems within the DIE	Major systems within the DIE	IPv4/6 products fail to match the need of the developed architecture.	6	9	7	Unknown/TBD	Possible: Because the expectation if that most of the hardware and software will be COTS and the ADO does not have complete control over the commercial suppliers.	Severe: Because parts of the DIE cannot be implemented at the required time causing loss of functionality and interoperability.	3	Possible	5	Severe	High	Risk must be treated as this risk has such wide-ranging impact.	Ensure that architecture developers fully understand the COTS roadmaps and the probability of suppliers meetings those roadmaps. Develop flexible architectures that can cope with varying implementations. Develop fall-back plans and investigate in-house solutions/patches using software solutions.
6	IPv6 Workshop (Group)	DIE	DIE	Schedule driven project delivery causes breakaway from the	7	19	20	Unknown/TBD	Possible: Because the schedules of the CDG and DMO	Major: Because this may result in a loss of	3	Possible	4	Major	High	Risk must be treated as this risk has	Determine set of metrics to monitor during course of implementation

#	Risk Author	Affected Components or Systems	Component or System Description	Description of Risk	Sources of Risk			Evaluation of Existing Controls	Likelihood of Risk	Consequences of the Risk	#	Likelihood Rating Value	#	Consequence Rating Value	Overall Rating (from DoD matrix)	Accept or Treat Risk? (with reasons why)	Treatment Strategies
				planned IPv6 implementation					are subject to external forces that the ADO does not have complete control over.	functionality and or interoperability within the DIE and between Allies.						such wide-ranging impact.	. Use results to apply changes to projects schedules to avoid breakaway.
7	IPv6 Workshop (Group)	DIE	DIE	Failure to manage schedules of the interdependent IP systems	7	19	20	Unknown/TBD	Possible: Because the schedules of the CDG and DMO are subject to external forces that the ADO does not have complete control over.	Major: Because this may result in a loss of functionality and or interoperability within the DIE and between Allies.	3	Possible	4	Major	High	Risk must be treated as this risk has such wide-ranging impact.	Governance mechanisms (IPv6PO) are designed to ensure that inter-dependant projects schedules can be managed. Metrics should be put in place to determine the extent of failure as soon as possible, treatment strategies could include strengthening the Governance measures, increasing budget and man-power.
8	IPv6 Workshop (Group)	DIE	DIE	Failure to manage technical standards between interdependent IP systems	7	19		Unknown/TBD	Possible: If the treatment strategies are not followed or fail.	Major: Because this may result in a loss of functionality and or interoperability within the DIE and between Allies.	3	Possible	4	Major	High	Risk must be treated as this risk has such wide-ranging impact.	Governance mechanisms (IPv6TO) are designed to ensure that technical standards between interdependent projects can be managed. A technical audit process should be put in place to determine the extent of non-compliance (standards failure) as soon as possible. Treatment strategies could include strengthening the Governance measures, increasing budget and man-power or determining the lowest cost method of re-aligning the standards.
9	IPv6 Workshop (Group)	DIE	DIE	Don't capture future IPv6 address space	9	6	7	Unknown/TBD	Possible: If the treatment strategies are not followed or fail	Moderate: Because the result may mean the ADO ends up with a non-	3	Possible	3	Moderate	Medium	Risk must be treated as this risk has such wide-	The IPv6 address plan should be regularly revisited to determine trends well

#	Risk Author	Affected Components or Systems	Component or System Description	Description of Risk	Sources of Risk			Evaluation of Existing Controls	Likelihood of Risk	Consequences of the Risk	#	Likelihood Rating Value	#	Consequence Rating Value	Overall Rating (from DoD matrix)	Accept or Treat Risk? (with reasons why)	Treatment Strategies
										contiguous address space and this affect routing performance.						-ranging impact.	ahead of time, so that solutions can be trailed on the IPv6 test-bed.
10	IPv6 Worksh op (Group)	DIE	DIE	Don't have skills/competencies to manage IPv6 transition	7	19	20	Unknown/TBD	Possible: Many of these skills will need to be supplied by organisation's external to the ADO.	Major: Because this may result in a loss of functionality and or interoperability within the DIE and between Allies.	3	Possible	4	Major	High	Risk must be treated as this risk has such wide-ranging impact.	Fund training of individuals to gain these skills. Co-operate with Allied (and other) agencies and embark upon a secondment program. Slow down the transition schedule to meet the reduced resourcing/skill level.
11	IPv6 Worksh op (Group)	DIE	DIE	IPv6 Transition Office not adequately resourced	7	19	20	Unknown/TBD	Possible: Because of funding restrictions or the inability to find these skills external to the ADO.	Major: Because this may result in a loss of functionality and or interoperability within the DIE and between Allies.	3	Possible	4	Major	High	Risk must be treated as this risk has such wide-ranging impact.	Increase funding and increase resourcing. Slow down the transition schedule to meet the reduced resourcing level.
12	IPv6 Worksh op (Group)	DIE	DIE	Fractured / poorly co-ord engineering processes and environments, e.g. test beds	7	19	20	Unknown/TBD	Possible: If the treatment strategies are not followed or fail.	Moderate: Because the test beds may not produce desired results for the planning process.	3	Possible	3	Moderate	Medium	Risk must be treated as this risk has such wide-ranging impact.	Governance mechanisms (IPv6TO) are designed to ensure that adequate engineering processes (inc test-bed environment) are created. A technical audit process (and evaluation process) should be put in place to determine the effectiveness of the developed processes. Treatment strategies could include strengthening the Governance measures, changing the processes, increasing budget and man-power.
13	IPv6 Worksh op (Group)	Affected DIE Applications	Affected DIE Applications	Cost of migrating the applications is significantly greater than planned.	7	6	13	Unknown/TBD	Possible: Because there are many applications within the DIE (not all	Major: Because may cause loss of funding for other parts of the	3	Possible	4	Major	High	Risk must be treated as this risk has such wide-	Increase funding. Delay migration of some non-critical applications (those not

#	Risk Author	Affected Components or Systems	Component or System Description	Description of Risk	Sources of Risk		Evaluation of Existing Controls	Likelihood of Risk	Consequences of the Risk	#	Likelihood Rating Value	#	Consequence Rating Value	Overall Rating (from DoD matrix)	Accept or Treat Risk? (with reasons why)	Treatment Strategies	
								COTS) and the estimation process necessarily will have a degree of error.	transition.						anging impact.	affecting interoperability) that can be isolated in enclaves of IPv4 for longer than planned. This delay could be permanent for the most expensive (to transition) applications.	
14	IPv6 Workshop (Group)	DIE	DIE	Independent (esp wrt to funding) stakeholder organisations don't comply with IPv6 policy/plan	7	19		Unknown/TBD	Possible: If the treatment strategies are not followed or fail.	Major: Because this may result in a loss of functionality and or interoperability within the DIE and between Allies.	3	Possible	4	Major	High	Risk must be treated as this risk has such wide-ranging impact.	It is assumed that all required ADO organisations will follow the IPv6 plan and the governance measures have been designed to achieve this goal. The governance measures will be weakest however for external organisations (e.g. Allies). The ADO should therefore extend its Communication s/Education program as far as possible to bring those stake-holders into the fold. Alternatively, measures (and fall-back plans) may need to be considered to alter the ADO plan.
15	IPv6 Workshop (Group)	External DIE interfaces	External DIE interfaces	IP services into external orgs may need to be maintained at IPv4	6	11		Unknown/TBD	Likely: Because we know that there are subject organisations who have not progressed very far down the IPv6 transition path.	Moderate: Because this will increase costs for the ADO and extend the transition period.	4	Likely	3	Moderate	Medium	Risk must be treated as this risk has such wide-ranging impact.	Although it is fully expected that some IP services will remain IPv4 for a long time into the future, there may be some which it is very desirable/nece ssary to switch to IPv6. The ADO could assist these organisations to make the transition more quickly by providing technical/managerial support, training and even funds.

#	Risk Author	Affected Components or Systems	Component or System Description	Description of Risk	Sources of Risk		Evaluation of Existing Controls	Likelihood of Risk	Consequences of the Risk	#	Likelihood Rating Value	#	Consequence Rating Value	Overall Rating (from DoD matrix)	Accept or Treat Risk? (with reasons why)	Treatment Strategies	
16	IPv6 Workshop (Group)	DIE	DIE	IPv6 plan not responsive to speed of development of COTS	7	6	19	Unknown/TBD	Possible: Because commercial pace of change is significantly faster than the non-commercial pace, this is just a plan and it cannot be perfect.	Moderate: Because this may create a lost opportunity for the ADO.	3	Possible	3	Moderate	Medium	Treat: Because it is better not to incur a lost opportunity.	Alter the plan as required to meet the actual pace of COTS development, this would be verified by testing products on the IPv6 test-bed.
17	IPv6 Workshop (Group)	DIE	DIE	Having an adequate range of IPv6 products on the approved products list (APL)	7	6	19	Unknown/TBD	Rare: Because we know there are already IPv6 products in the market place, the only restriction is to go through the ADO processes to get them on the APL.	Minor: Because it is assumed that there will be other products on the APL that can do the job, the only impact is that you may not end up with the optimum implementation.	1	Rare	2	Minor	Low	Accept: This is a nice to have only.	None required.
18	IPv6 Workshop (Group)	DIE	DIE	Risk management context not defined.	7		19	Unknown/TBD	Rare: Because this is easily solved and largely an ADO management issue.	Minor: Because the effect should only generate less finely tuned or out-of-scope risks.	1	Rare	2	Minor	Low	Treat: Because effort (cost/schedule) could be wasted treating non-risks.	Undertake during the early part of the detailed-planning phase, a study to determine this context in detail.
19	IPv6 Workshop (Group)	DIE	DIE	ISB don't have a rigorous enough process to detect IPv6 enabled equipment that is connected to the network and may cause problems.	7	6	19	Unknown/TBD	Rare: Because we know this is technically possible and solvable so this is largely a management issue.	Moderate: Because the effect could compromise security or cause network performance problems	1	Rare	3	Moderate	Low	Treat: Because this should be straight forward to achieve.	Using the resources of the IPv6TO to trial a better process on the IPv6 test bed and work with ISB to improve the situation.
20	IPv6 Panel (John Pennington)	Affected DIE Applications	Affected DIE Applications	Application transition turns out to be more difficult than expected	7	6	18	Unknown/TBD	Possible: Because the work to evaluate applications for transition is TBC and we know that there are non COTS applications that could be expensive to transition.	Moderate: Would increase costs and may affect budget for other areas of the transition.	3	Possible	3	Moderate	Medium	Treat: Because cost and schedule may be involved.	Assuming that the budget and schedule is soaked up in transitioning less applications, the solution may be to accept that more applications live on in IPv4 for longer. Alternatively more budget is sought to transition the remaining applications and/or a more rigorous process is undertaken to either find ways that the applications

#	Risk Author	Affected Components or Systems	Component or System Description	Description of Risk	Sources of Risk			Evaluation of Existing Controls	Likelihood of Risk	Consequences of the Risk	#	Likelihood Rating Value	#	Consequence Rating Value	Overall Rating (from DoD matrix)	Accept or Treat Risk? (with reasons why)	Treatment Strategies
																	can be removed from service or replaced by other processes or applications.
21	IPv6 Panel (John Pennington)	DIE	DIE	Tactical comms equipment is not available to support IPv6 in target timescale	14	6	20	Unknown/TBD	Likely: Because we know that JP2072 is still in the early stages of development and the JTRS program is in delay.	Major: Because we will need to maintain IPv4 for operational systems and will lose advantages of IPv6.	4	Likely	4	Major	High	Treat: Because the tactical space is crucial to the ADO ability to carry out operations.	Either increase funding to pull-forward tactical equipment availability, or find an interim capability that can be delivered earlier, or support legacy systems for longer.
22	IPv6 Panel (John Pennington)	DIE	DIE	Network design needs nested tunnels (e.g. for cryptos, routing encapsulation) but MTU limits are breached	6	18	11	Unknown/TBD	Possible: Because this is a technical issue and the solution is TBD.	Minor: Because some network connections fail, or expensive work-arounds needed.	3	Possible	2	Minor	Medium	Accept: Because work-around is acceptable.	Redesign the network to avoid this situation. There may be some network equipment where the IP layer is implemented in software and the manufacturer "may" be able to provide a work-around, however this is not recommended.
23	IPv6 Panel (John Pennington)	DIE	DIE	Evaluated firewall, CND and crypto products for IPv6 not available in time	6	14	18	Unknown/TBD	Possible: Because security products tend to take longer to be made available than general purpose commercial COTS infrastructure.	Moderate: Because target dates not met, cost to reschedule projects	3	Possible	3	Moderate	Medium	Treat: Because of cost and schedule impacts.	Apply more resources to the accreditation process if this is the bottleneck. Otherwise if this is a COTS availability problem then either delay the role out or consider finding ways to assist with the suppliers meeting the ADO's need.
24	IPv6 Panel (John Pennington)	DIE	DIE	PKI solution not available to support IPsec, either in ADO, or to allies	6	14		Unknown/TBD	Possible: Because commercial security products not under complete control of ADO.	Minor: Because greater security capability not available.	3	Possible	2	Minor	Medium	Accept: Because security products are already in place.	None required.

#	Risk Author	Affected Components or Systems	Component or System Description	Description of Risk	Sources of Risk		Evaluation of Existing Controls	Likelihood of Risk	Consequences of the Risk	#	Likelihood Rating Value	#	Consequence Rating Value	Overall Rating (from DoD matrix)	Accept or Treat Risk? (with reasons why)	Treatment Strategies
25	IPv6 Panel (John Pennington)	DIE	DIE	Windows Active directory does not migrate to IPv6 in time	1	6	14 Unknown/TBD	Possible: Because commercial products not under complete control of ADO.	Severe: Because the DIE transition must be delayed	2	Unlikely	5	Severe	High	Treat.	If the ADO uses Microsoft Active Directory (AD) widely then a delayed IPv6 transition will have to be accepted, except for specific systems where it is essential (Allies) however the application gateway approach may be a less cost lower risk solution. If the use of AD is not widespread then these systems could be enclaved and maintained as IPv4 until Microsoft delivers support.
26	IPv6 Panel	Affected DIE hardware	Affected DIE hardware	Cost of migrating the hardware is significantly greater than planned.	14	13	6 Unknown/TBD	Rare: Because it is expected that only general purpose (e.g. printers etc) peripherals will be affected.	Insignificant: Because the solution is to continue to support IPv4 in the DIE and this is planned for some time.	1	Rare	1	Insignificant	Low	Accept: Because the work-around has already been identified in the plan.	None required.

Source of Risk		
0		
1	Commercial and legal relationships	Between the organization and other organizations, e.g. suppliers, subcontractors, lessees.
2	Economic circumstances	Of the organization, country, internationally, as well as factors contributing to those circumstances e.g. exchange rates.
3	Human behaviour	Of both those involved and those not involved in the organization.
4	Natural events	
5	Political circumstances	Including legislative changes and factors which may influence other sources of risk.
6	Technology and technical issues	Both internal and external to the organization.
7	Management activities and controls	
8	Individual activities	
9	Materiel System requirements	Materiel System requirements, as defined in the OCD and FPS (noting that inadequate requirements are identified as the No 1 cause for project failure);
10	Operating environment	Operating environment (i.e. how similar is the operating environment for which equipment was designed with the envisaged operating environment?);
11	Interfaces	
12	Software development and management	Software development and management, including software support;
13	Degree of development required for the system	
14	Maturity of technology required	
15	Specialty engineering areas	Specialty engineering areas, such as growth and obsolescence, safety, security, electromagnetic environmental effects, human factors, and radio-frequency spectrum management;
16	Government Furnished Material	Government Furnished Material (GFM), which includes Government Furnished Equipment (GFE), Government Furnished Data (GFD) (i.e. warranted data), and Government Furnished Information (GFI);
17	Integrated Logistic Support	Integrated Logistic Support (ILS) issues, including: Support System requirements, support contract requirements, linkages between the acquisition and support contracts, and costing and resourcing the envisaged support arrangements;
18	Transition from an existing Materiel System	Transition, particularly the transition from an existing Materiel System (or part thereof) to a new Materiel System (while maintaining capability);
19	ADO project offices	ADO project offices, particularly with respect to the right balance of personnel numbers, skills and experience; and
20	Defence contractors	Defence contractors, particularly with respect to capability to undertake the required work (i.e. process maturity and the right balance of personnel numbers, skills and experience).

ANNEX D DCP Project Summary

Project Number	Title	Relevance
DEF 7013	Joint Intelligence Support System (JISS)	Further development of the JISS for support of the Australian Defence intelligence community.
AIR 5276 Ph 6	Data links for AP3-C Orion aircraft.	Upgrade aircraft communications suite and data links. Note: Currently use Link-11.
AIR 6000	Joint Strike Fighter (JSF).	Comms/Radios will come as part of this platform acquisition
AIR 7000	Multi-mission Maritime Aircraft (MMA). AP3-C replacement.	Comms/Radios will come as part of this platform acquisition
AIR 9000	Helicopters.	Comms/Radios will come as part of this platform acquisition.
JP 2008	Military Satellite Communications	Expanded capability including use of Optus/Singtel C1 satellite.
JP 2030	Joint Command Support Environment	Consolidating existing Command Support Systems into a single environment.
JP 2047	Defence Wide-Area Communications Network	Multi-phase project with ISDs between 2005 and 2014. Providing enhanced encryption services, enhanced protocols transmission and switching equipment and providing guidance of on-going development.
JP 2068 ⁷²	DNOC –Defence Network Management System and Computer Network Defence	Improving management, monitoring, security and visibility of the DIE.
JP 2069	High Grade Cryptographic Equipment (HGCE).	Replacement HGCE.
JP 2072	Battlespace Communications System (Land)	Replacing the Army's CNR and Tactical Trunk Communications with an advanced communications system.
JP 2089	Tactical Information Exchange Domain (TIED) (Data Links)	Delivering Link-16 and VMF on Ships and Planes/Helicopters and associated land-based platforms.
JP 2090	Combined Information Environment	Establish permanent "information" connectivity between ADF and key Allied Command and Control networks and systems to support future Coalition operations.
LAND 75	Battlefield Communications Support System (BCSS)	Role out of BCSS below Brigade level.
LAND 125	Soldier Combat System	Acquire advanced capabilities for the combat soldier.

⁷² It was advised during the IPv6 Workshop that Phase 2A of JP2068 has been cancelled and that JP2047 will provide NMS functionality enhancements.

SEA 1442	Maritime Communications and Information Management Architecture Modernisation	Introduction of Maritime Tactical Wide Area Network and IP Networking to a range of RAN vessels.
SEA 4000	Airwarfare Destroyer	Comms/Radios will come as part of this platform acquisition

ANNEX E IPv4

IPv4 Address Space Exhaustion

One of the potential consequences of failing to transition from IPv4 to IPv6 may be the exhaustion of IPv4 addresses. Figure 15 plots the allocation of IPv4 addresses against time and shows that prior to 1995 addresses were being allocated at a steep linear rate, these were mostly Class B⁷³ allocations. Since 1995 a CIDR methodology has been used to allocate addresses and Figure 27 also plots a prediction (green line out to 2015) of address allocation using an exponential model starting in 1995.

Using this exponential model the pool of un-allocated IPv4 address will be exhausted by February 2014⁷⁴.




Figure 35 IPv4 IANA Allocations - Projection using Exponential Growth Model⁷⁵

It should be noted that the above does not necessarily relate to the IPv4 address usage within the ADO.

⁷³ A Class B address range will support up to 65,534 Hosts.

⁷⁴ Source <http://bgp.potaroo.net/ipv4/> , this chart is automatically updated each day.

⁷⁵ Source <http://bgp.potaroo.net/ipv4/> , this chart is automatically updated each day.

ANNEX F CIOG Organisations Chart

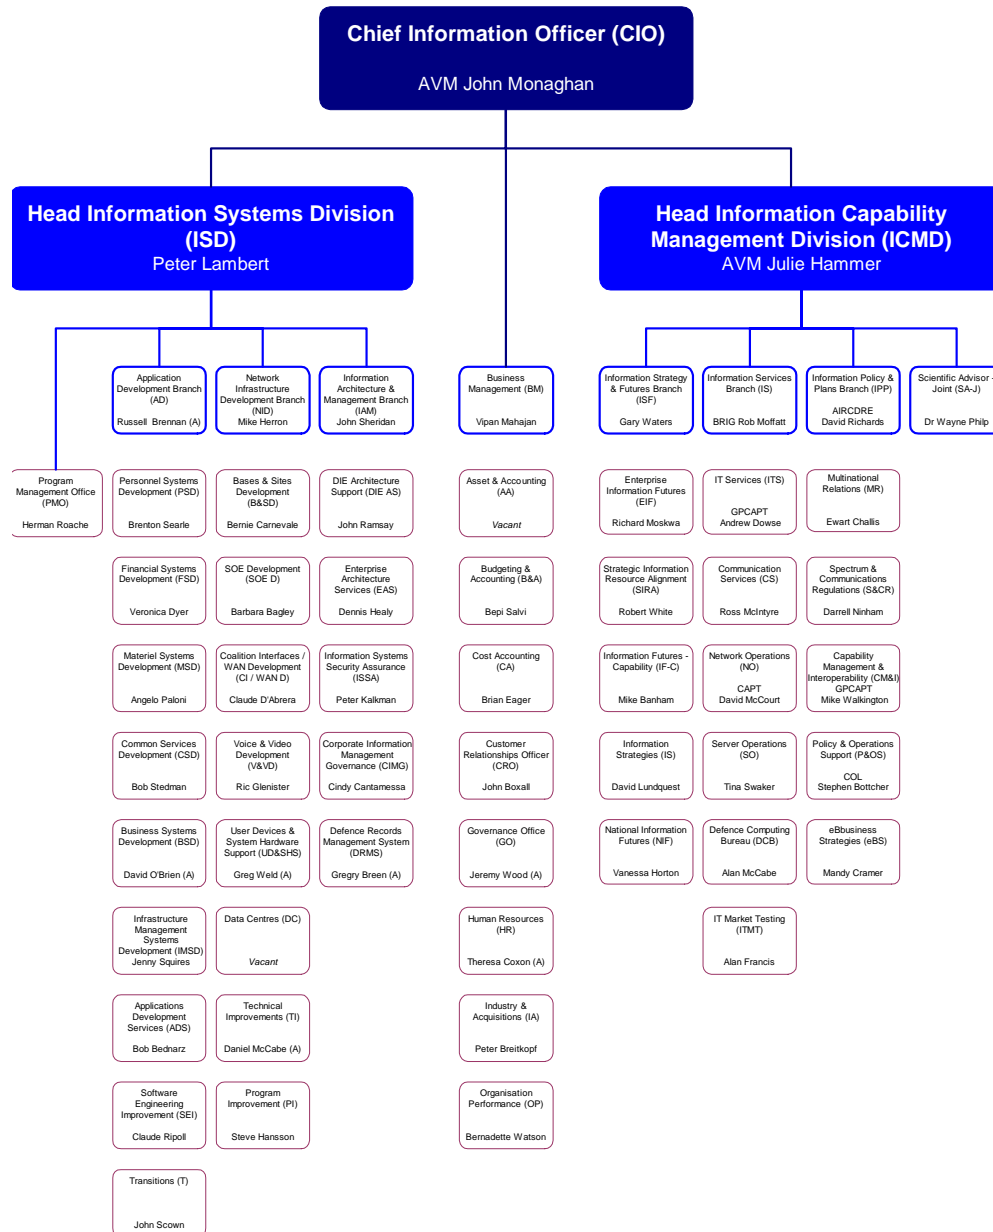


Figure 36 CIOG Organisational Structure

ANNEX G Mobile IP

Mobility in IPv6




Figure 37 Edge Mobility

Mobile IPv6 is designed to support individual roaming mobile hosts. The aim of MIPv6 is to maintain reachability to a node as it moves to various points in a network. The mobile node can be reached via a constant “home address” when it is on a different network. Active sessions can be maintained as the node moves from one network to another.

Mobile IPv6 has three main components: the mobile node (MN), the home agent (HA), and the correspondent node (CN). The way Mobile IPv6 works is as follows. The mobile node registers with a specific home agent. When the MN moves to a new network (presumably by connecting to some sort of access point), it must detect that it has moved to a new network, and obtain a new IP address. While a new address can be obtained via DHCPv6 after being triggered by some sort of movement detection process, the more common method uses router advertisements (RADV) to detect movement and automatically assign a new address using stateless auto-configuration. Once this new address is obtained, the mobile node sends a binding update (BU) to the HA and any correspondent nodes it is currently communicating with, notifying them of its new care-of address (CoA).

There are two possible modes of communication between the mobile node and a correspondent node. The first mode, bi-directional tunnelling, does not require MIPv6 support at the correspondent node. In this mode, the home agent intercepts all packets destined for the MN using proxy neighbour discovery, and tunnels them to the MN. Packets that the MN sends to the CN are also tunnelled back through the home agent. The second mode of communication, route optimisation, requires the correspondent node to have Mobile IPv6 functionality. This process starts out the same as the bi-directional tunnelling mode, with the HA intercepting packets destined for the MN, and tunnelling them to the MN's Care-of Address (CoA). With route optimisation, the MN then informs the CN about its CoA, and the CN and MN can then communicate directly, without the aid of the HA. As long as a session is active, the MN needs to send a binding update (BU) to the CN when it moves to a new network, so that they may continue direct communications.

General MIPv6 Benefits

Here are the primary similarities and differences between MIPv6 and MIPv4:

- **Foreign agent.** Both standards rely on a home agent and a mobile node, but MIPv6 does not define a foreign agent to issue a care-of address (CoA), since routable address constraints are not an issue in IPv6 networks. Instead, MIPv6 derives the CoA directly from auto-configuration schemes. This approach enables the mobile node to operate in any location without requiring special support from the local router.
- **Route optimisation.** MIPv6 enables direct-packet routing between the mobile node and corresponding nodes located on an IPv6 network. When the mobile node moves into a foreign network, it obtains a new CoA and reports this to its home agent. The home agent intercepts all packets destined for the mobile node and tunnels them to its registered CoA. In a MIPv4 scenario, a corresponding node's traffic must pass through the home agent, but MIPv6 route optimisation allows the mobile node to send binding updates to an IPv6-based corresponding node. The corresponding node caches the current CoA and then sends packets directly to the mobile node. This is an optional procedure for MIPv4 that requires special options to be enabled on each corresponding node, and is rarely implemented or used.
- **Security.** MIPv4 and MIPv6 will often be used with a VPN (virtual private network) solution for data security when the user is roaming into networks outside the corporate firewall. Both protocols will in theory allow the use of a v4 IPsec (Internet protocol security) VPN solution, providing in the case of the MIPv6 client that the IPv6 protocol stack includes a 6-to-4 function. In addition, the MIPv6 client allows the use of a v6 IPsec VPN solution.
- **Home agent address discovery.** Using the IPv6 anycast feature, the mobile node can send a binding update to the home agent anycast address. The mobile node will get only one response from one home agent even if several are present on the network. This is an efficient way of keeping track of multiple home agents, which may be required in many networks for redundancy or scalability.

MIPv6 status

MIPv6 has mature IETF standards-track specifications for its core functionality, as well as for the added ability to use IP Security (IPSec) to encrypt signalling between the MN and the HA. There are a few reasonably mature MIPv6 implementations available covering the Linux, BSD, CISCO IOS, and Windows operating systems, as well as simulation environments.

There is still significant evolving research being done in the area of MIPv6. Emerging enhancements and modifications, such as Hierarchical MIPv6 (HMIPv6) may help improve the performance and scalability of the protocol.

General MIPv6 Issues

MIPv6 is only necessary for mobile end systems which require a stable IP address for identification or to maintain in-progress sessions while roaming between networks. If these conditions need not be met, and it is acceptable to obtain a new address and restart current sessions, then the combination of DHCP and dynamic DNS, as one possible example, may be sufficient to meet mobility criteria, and MIPv6 maybe unnecessary.

The main MIPv6 specification includes a mechanism for Dynamic Home Agent Address Discovery (DHAAD), which can be used for avoiding a manual configuration of the Mobile Node with the Home Agent's address. However, the current mechanism that allows a Mobile Node to detect the prefix of its home network when attached to a visited network, requires additional operational administration. This must currently be done manually, though there is work underway to address this aspect in an automatic way using the Authentication, Authorization, and Accounting (AAA) infrastructure, and new methods to identify new link prefixes from work on Detecting Network Attachment (DNA) which is work in progress.

If tunnelling is used between the MIPv6 Home Agent and Correspondent the standard tunnelling overhead for any protocol will exist, but this can be avoided using the MIPv6 route optimisation. Additionally, fast handoff of a mobile node has significant limitations due to the required local interface protocol standards.

Network Mobility

Network Mobility (NEMO) is essentially an extension to Mobile IPv6. NEMO is designed to apply to entire networks in motion, rather than just individual nodes in motion. It is still an area of work in progress within the IETF.

2. GES SOA Standard Review

Memorandum for Director Defense Information System Agency

From: Executive Director, W2COG Research Initiative

Ref: (a) DISA Proposed Standards for Implementing GIG Enterprise Services

Encl: (1) W2COG Institute Response to Reference (a)
(2) Prof Rick Hayes-Roth, Naval Postgraduate School, Response to Reference (a)
(3) Object Management Group (OMG) Response to Reference (a)
(4) Proposed Statement of Work for Phase II Evaluation

Subj: W2COG Review of Proposed Standards for Implementing GIG Enterprise Services

1. Thank you for the opportunity to review ref (a). Appendix (1) is a good faith W2COG Institute quick look, pending opportunity for a more thorough effort. It provides a collaborative industrial perspective that endorses your use of non-proprietary standards to achieve an interoperable, affordable, and leading (not bleeding) edge Commercial-Off-the-Shelf (COTS)-based solution that avoids interoperability problems and vendor lock-in; agrees that your proposed list is a viable beginning; and opines that industry will marshal behind DISA in support of a non-proprietary framework. The Institute suggests an expansion of the framework with special attention to building distributed, high performance, highly secure and highly reliable systems using fundamental SOA concepts. It comments briefly on the issue of test and validation for SOA-based enterprise services, noting that it makes little practical sense to be able to implement a SOA-based solution in a matter of months only to be forced to spend years in testing it and validating it prior to deployment. Finally it strongly suggests follow-on review by carefully selected expert team.

2. Encl (2) provides an opinion that frames the essential basis of the W2COG value proposition. Its author is a renowned expert in industrial software and architecture and is one of the founders of the W2COG project. His point is that “(choosing) standards (before demonstrating enhanced capability) puts the cart before the horse. Instead, we should focus on high-value transactions that achieve information superiority, and then choose tools, including, but not chiefly, standards that reduce *time to value*.”

3. Appendix (3), provided by the president of a very large and respected software standards body, suggests that software standards for architectural design, e.g. UML(2), should be included along with standards for architectural implementation..


4. Appendix (1) demonstrates how quickly and cost-effectively government can apply the not-for-profit umbrella of the W2COG to leverage on going commercial investment in software architectural development. Accordingly, appendix (4) suggests the parameters of a phase II review of reference (a).

5. I personally endorse all the above and add further comment. One could argue that identification and implementation of COTS software standards will mostly “just happen” agnostic of DoD’s efforts to mandate its list of standards. On the other hand, DoD must be vitally concerned about its process for dynamically selecting and applying the situationally appropriate standards to field the capability it desires. It seems to me that the NCES technical references cited in reference (a) address the former more than the latter. To close this gap we might look at past lessons learned. For example, we could take an approach similar to the Defense Acquisition University “breathalyzer” test for software development from the days of MIL-STD-2167. Software developers were required to answer a short list of tough questions keyed to acquisition imperatives before fielding capability. If you agree that DISA’s technical reference material should facilitate creation of a dynamic SOA standard implementation process, you might take a fresh look at the “breathalyzer’s” list of ten questions and apply them to design technical reference for dynamic NCES software implementation process. If you are interested in pursuing this line of thought, we have some ideas and can role them into the phase II evaluation.

C. R. Gunderson

World Wide Consortium for the Grid (W2COG)

1895 Preston White Dr
Reston VA 20191
(703) 262 5332
www.w2cog.org



January 20, 2006 *via email transmission to john.rosenbaum@disa.mil*

Defense Information Systems Agency
Attn: Dr. John Rosenbaum
PO Box 4502
Arlington, VA 22204

SUBJECT: Proposed Standards for Implementing GIG Enterprise Services

Dear Dr. Rosenbaum:

Please accept this input as the W2COG Institute preliminary response to LTG Charles E. Croom, Jr. Memorandum of 27 December 2005, subject as above.

W2COG Institute and its affiliated companies are pleased to submit this response as part of our continuing research program directed toward accelerating the development and availability of tools to support secure, net-centric operations.

We share with you a congruence of goals centered on four mutually reinforcing characteristics for enterprise service standards:

- **Non-Proprietary.** Enterprise services built using a Service Oriented Architecture (SOA) must guarantee interoperability to achieve information sharing in a net-centric environment.
- **Distributed.** SOA-based services must satisfy the needs of a Global Information Grid and operate seamlessly when mobile or forward deployed.
- **High Performance.** SOA-based services must support high transaction rates and large volumes of enterprise data, and still exhibit low latency.
- **Secure and Reliable.** SOA-based services must be secure and ensure high rates of availability under adverse conditions.

Sincerely,



Aaron Budgor

Executive Summary

W2COG Institute concurs with comment to the proposed list of standards for implementing Global Information Grid (GIG) enterprise services. We endorse the use of non-proprietary standards to achieve an interoperable, affordable, and leading (not bleeding) edge Commercial-Off-the-Shelf (COTS)-based solution that avoids interoperability problems and vendor lock-in.

We believe that the list of standards proposed by the Net-Centric Enterprise Services (NCES) Program Office is a viable beginning to achieving the goal of a non-proprietary instantiation of Service Oriented Architectures (SOA). We further believe that industry will marshal behind DISA in support of a non-proprietary framework.

In this paper we will build upon the work of the NCES Program Office to briefly suggest an expansion of the framework to better support the net-centric environment with special attention to building distributed, high performance, highly secure and highly reliable systems using fundamental SOA concepts.

We comment briefly on the issue of test and validation for SOA-based enterprise services, noting that it makes little practical sense to be able to implement a SOA-based solution in a matter of months only to be forced to spend years in testing it and validating it prior to deployment.

To ensure widespread industry and defense stakeholder support for an expanded list of standards, we also explicitly suggest a Phase II comment period, not to exceed 60 days, during which a more complete analysis and architectural review could occur. We offer industry support from W2COG Institute and its affiliates to participate in such a review.

Success Criteria

In reviewing the list of standards proposed by the NCES Program Office, the W2COG Institute applied the following success criteria:

- **Industry Support.** Will private industry support the proposed standards? The premise of the NCES effort is that core enterprise services of the Department of Defense can be implemented using proven commercial technologies from the Internet and electronic business. Because of its network of members and affiliates, W2COG Institute is in a prime position to judge whether the proposed standards will garner sufficient industry support to realize this goal.
- **Sufficiency and Practicality.** Can high performance systems with excellent security and reliability be built with these standards? We asked experts who have actually built multiple high performance SOA-based systems for private industry to comment on perceived gaps and overlaps.
- **Sustainability.** Does the proposed list of standards lay a solid foundation for sustainable, manageable effort? Or, will it die of bureaucracy and overhead? We again looked at commercial best practice to compare, but we also critically examined the theory of SOA to determine potential strengths and weaknesses.
- **Testability.** Can a SOA-based enterprise service which is based on these standards be efficiently tested and validated? We examined the standards carefully to determine if technological solutions or industry best practices to deal with the testing issue could be supported.

Endorsement of XML Base Protocol Standards

To level-set our understanding of a SOA-based enterprise service, we begin by providing a graphical depiction of what we mean by a Service Oriented Architecture. Briefly, a services approach to system design, implementation and provisioning is based on the following definitions:

Service – a coarse-grained (i.e. “moderately large”) business or technology capability unit with well defined interface boundaries that interacts with end users and other services through industry standard, message-based protocols.

SOA – Service Oriented Architecture – technical architectures based on event-driven collections of loosely-coupled software components which implement services.

SOE – Service Oriented Enterprise - An organization which uses the concept of a service to optimize its enterprise architecture. A “service” orientation implies a layered architecture with support for a business processes layer, a service oriented application architecture (SOA) layer, a service oriented infrastructure (SOI) layer, and a service management layer.

SOI – Service Oriented Infrastructure – a virtualized “landing zone” for SOA solutions in which hardware, storage, security and network resources are virtualized and managed as a utility.

As part of establishing this definition we explicitly accept the standards for XML base protocols as essential elements in any mandatory list of SOA implementation standards. The XML base protocols are SOAP, WSDL and UDDI. We note that these standards are shown in the NCES Program Office suggested list. The relationship and interaction among the XML base protocols are shown in the following diagram:

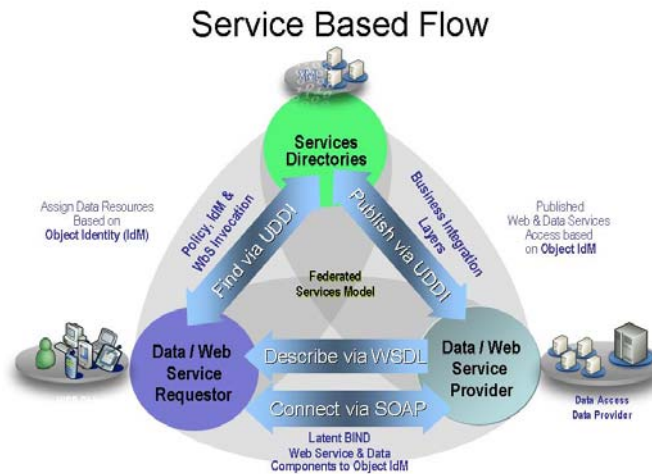


Figure 1 -- XML Base Protocols

Benefits and Caveats of a Standards-Based Approach

The W2COG Institute agrees with your desire to use standards to achieve an interoperable, affordable COTS-based solution that avoids individual vendor lock-in. Standards can define a common conceptual taxonomy for describing the problem space and encourage the use of modular, and thus re-usable and upgradeable solutions.

Unfortunately, dedication to standards alone does not guarantee success. Standards may be in flux, may be overly cumbersome, may be implemented differently by different vendors, and do not assure optimal performance, lowest cost of implementation, or efficient manageability. Consider the following extremes:

- Several of the web services standards listed by the NCES Program Office have been changed frequently by their approving bodies in the last two years. One best practice is to build an architectural layer of reusable implementation capabilities below the WS-* standard in order to more rapidly accommodate specification changes in the standard itself.
- Several of the standards proposed remain incomplete or unapproved by their authoritative body, or they remain unchanged when in fact they desperately need updates. The reasons for this are varied. Standards bodies do not always achieve perceived improvements in successive versions, particularly when a standard becomes more complex. In addition, once a critical mass of vendors has adopted a standard, there is often little business advantage to keeping up with changes.

In relation to the issue of updates and ever changing standards, we would like to call your particular attention to the following considerations:

- ebXML Registry is now a mature OASIS Standard and incorporates all the important trading-partner interaction standards such as OASIS Collaboration Protocol Profile Agreement (CPPA). It also enables the discovery of users and SOA Artifacts such as WSDLs and WS-Policy files. It should be listed alongside UDDI Registry.
- WS-Security has just gone to vote on Version 1.1. It should pass by end of this month at the latest. The standards should be updated to include support for this standard and version.
- While not currently being adopted for SOA Security functions, XrML is being evaluated for use to provide Content and Services rights management. It is also being looked at as a potential method of applying Policy rules to the use of Services and the data being served.
- The use of SOAP 1.1 should be re-evaluated relative to SOAP version 1.2. Technically, SOAP 1.1 is not a standard, in that it was superseded by SOAP 1.2 before it could be approved.
- We take cautionary exception to WebDAV, and suggest it be reconsidered. We believe it is relatively proprietary and presents a security vulnerability.

Thus, while we support the use of standards in designing and implementing the NCES architecture, our approach must first be based on meeting the desired results of achieving mission effectiveness: secure, deployable, reliable, and manageable solutions that meet service level agreements and maximize benefit versus cost.

Proposed Requirements Taxonomy

In order to examine the list of standards for gaps and overlaps, we propose the following taxonomy of functional requirements for SOA standards:

Access:	Authorization:	Orchestration:	Processing:
Access Ctrl	Authorization Services	Transformation	Data Adaptors
Reliable Asynch Communication	Assertion Mapping	Translation	WS Adaptors
Virtualized IP Layer 2 Security	SLA Mappings	Semantic Lib Normalization	Native Adaptors
XML FW/Inspection	LDAP/DS Integration	Biz Process – Rules Execution	Intelligent Routing
Access & Authorization Svcs		Composite Svcs Execution	Publishing/ Submission
Edge Data Persistence		Fine Grain Svcs Execution	Store/Forward
Edge WS Interactions		MSG Transformation	Events
		UDDI Registry	

Figure 2 -- SOA Technical Domains

Based on our experience, we feel that every requirements area in this taxonomy should be covered by one or more standards. We say one or more, because the standards listed in the NCES proposed list are not all at the same level of abstraction. Some of the more abstract standards are implemented using the more basic standards. In order to resolve the apparent inconsistencies, we will map the above requirements taxonomy to a “protocol stack” which is a logical framework showing the relationships among top level protocols.

Our logical framework is shown in the next diagram:

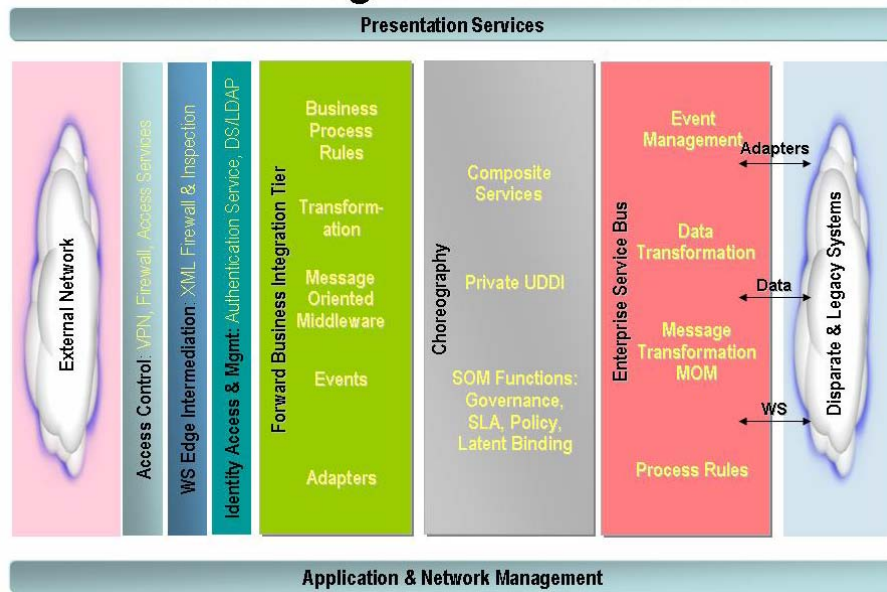


Figure 3 -- Logical Framework for SOA Standards

Then, while keeping in mind that the XML base protocols are ubiquitous and therefore assumed to be omnipresent in the logical framework, we can overlay the standards from the NCES proposed list, **and augment it with additional standards** that we feel are necessary to completely cover the requirements graph. The result is shown in the following diagram:

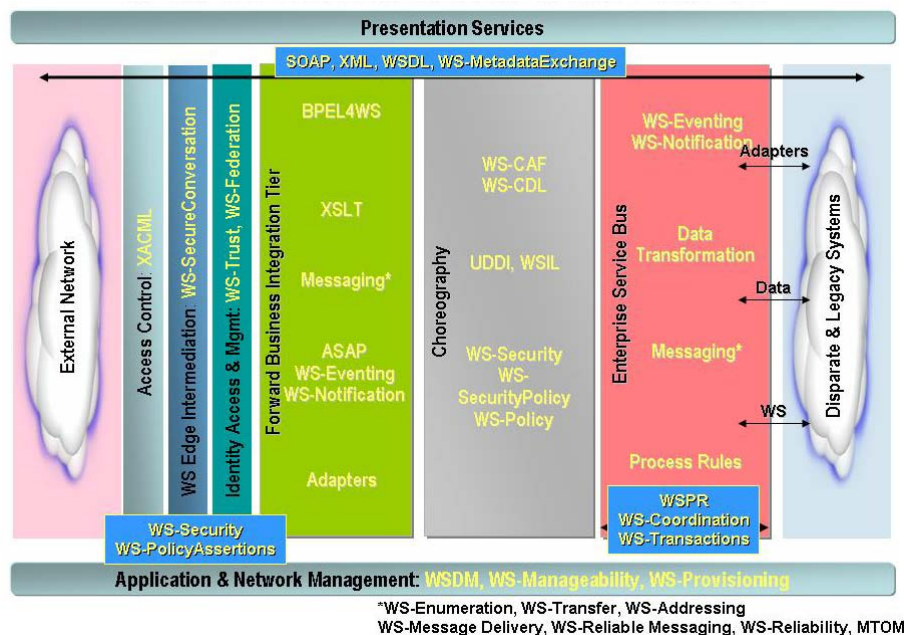


Figure 4 -- SOA Standards Overlay

AP Analysis

By comparing the original NCES proposed list to the coverage overlay from the requirements depicted in the logical framework, we can inventory the standards needed to complete the set. The following table outlines this gap analysis more precisely:

With the logical framework shown above to provide a visual context for related standards, this table directly responds to your list of proposed standards. It lists the architectural components of a SOA, related

standards, and the proposed standards you listed. When shown in red, we believe the standard is either an implicitly included implementation sub-standard, or it is a tool or technique and not a standard, *per se*.

Architectural Component	Related Standards	Proposed NCES Standard
Management		
○ Distributed Management	○ WSDM, WS-Manageability	
○ Provisioning	○ WS-Provisioning	
Security		
○ Security	○ WS-Security	○ WS-Security ○ XML-Signature ○ XML-Encryption
○ Security Policy	○ WS-SecurityPolicy	
○ Secure Conversation	○ WS-SecureConversation	
○ Trusted Message	○ WS-Trust	
○ Federated Identity	○ WS-Federation	○ SAML ○ XACML
Portal and Presentation	○ WSRP	CSS ○ WEBDAV ○ XSLT transformation
Transactions & Business Process		
○ Asynchronous Services*	○ ASAP	
○ Transaction	○ WS-Transactions, WS-Coordination, WS-CAF	
○ Orchestration	○ BPEL4WS, WS-CDL	
Messaging		
○ Events and Notification	○ WS-Eventing, WS-Notification	
○ Multiple Message sessions	○ WS-Enumeration, WS-Transfer	
○ Routing / Addressing	○ WS-Addressing, WS-MessageDelivery	○ SOAP
○ Reliable Messaging	○ WS-ReliableMessaging, WS-Reliability	
○ Messaging Packaging	○ SOAP, MTOM	
Metadata		
○ Publication and Discovery	○ UDDI, WSIL	○ UDDI ○ ebXML Registry
○ Policy	○ WS-Policy, WS-PolicyAssertions	
○ Base Service and Message Description	○ WSDL	○ WSDL
○ Metadata Retrieval	○ WS-MetadataExchange	○ WS-I interoperability

Table 1-- Gap Analysis

It is our recommendation that the standards listed in the center column be considered, when not already listed, for inclusion in the NCES Program Office proposed list of standards.

Reference Architecture

To further clarify the value of including the more complete list of standards, we offer the following reference architecture:

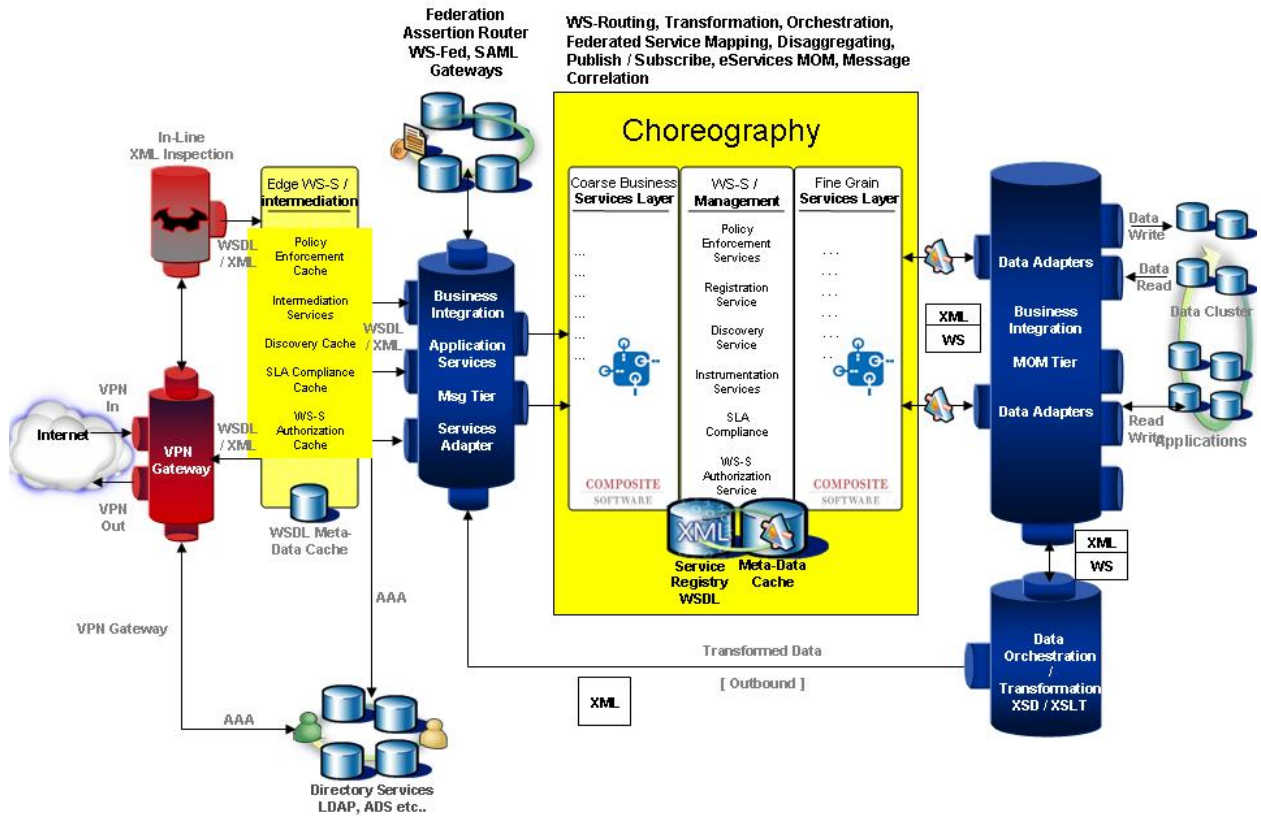


Figure 5 -- Services Reference Architecture

This reference architecture was prepared with the assistance of W2COG Institute members and other industry partners including Accenture and CSC. We believe it is the best available depiction of standards-based, non-proprietary support for SOA-based design.

Testing and Validation

As an additional comment on the proposed list of standards, we suggest that special caution is required in the area of testing and validation. We suggest mimicking commercial best practice in order to shorten the testing and validation period prior to deployment. This is vitally needed because it makes little sense to shorten the development parts of the life-cycle and still not reap the rewards of COTS procurement because of excessively long testing periods.

Our industry membership reports success with the following methods:

- A technique known as “test-centered design” is particularly appropriate for use with SOA because SOA includes a rigorous definition of the interface specification using the XML base protocols.
- Automated tools for the generation and execution of unit and integration tests are used to remove human factors and associated delays in tests prior to full system integration.
- Where feasible from a business perspective, rapid roll-back procedures and the concept of a “stateless enterprise” are used to rapidly provision new systems in the actual production environment with the cautionary capability of rapid return to the previous system in event of catastrophic failure.
- Continuous process improvement techniques such as “Six Sigma” are used to measure and lower defect rates in software production and system provisioning.

W2COG Institute would welcome the opportunity to define these and other methods from industry in more depth in order to provide DISA with the requisite best practices in testing and validation.

Call To Action

W2COG Institute believes that the NCES Program Office proposed list of SOA standards should be adopted.

We would also suggest that the standards gap analysis we have shown here may be of value in extending the list to a more complete and uniform list, and that our reference architecture provides a more meaningful way to depict and explain the interactions among the proposed standards.

In order to socialize these proposed enhancements to the standards list, we propose a period of continued education, research and debate into the viability of the proposed additions and their respective versions as approved by their respective standards bodies and endorsed through industry use. Particular attention should be paid to emerging standards in the following areas:

- Security.
- Infrastructure.
- Distributed Management.
- Business Process Coordination.

MEMORANDUM

TO: Chris Gunderson, W2COG

SUBJECT: *Proposed Standards for Implementing GIG Enterprise Services*

FROM: Rick Hayes-Roth

DATE: January 12, 2006

I've read the DISA memo on proposed standards. I'm familiar with the proposed standards. While CTO/Software at HP, I had the opportunity to fund and oversee HP's participation in several of these standards, as well as several others that might also have been usefully added to this list. In my current position as Professor of Information Sciences at the Naval Postgraduate School, I regularly analyze technology strategies and policies. I understand the intent and appreciate the value of standards. Nevertheless, I have mixed reactions to this proposal for these standards. In brief, I'm concerned that someone believes that a standards profile represents a significant, stable building block or, worse, that these standards can be a principal foundation for delivering significant value, quickly, to joint and coalition forces. I think standards put the cart before the horse. Instead, we should focus on high-value transactions that achieve information superiority, and then choose tools (including, but not chiefly, standards) that reduce *time to value*. Because such an opinion may sound heretical, in the section after next, I'll expose my reasoning more fully. However, before explaining my criticism, I want to suggest a constructive approach for avoiding the perceived risks, leveraging appropriate standards, and achieving more significant results more quickly.

Recommended approach:

Let's put the horse before the cart: To achieve information superiority, the goal of the GIG/NCES, the Pareto-optimal approach is to implement the highest-value mission "threads" or "transactions" first, accumulating lessons, technology components, patterns, and architectural frameworks along the way. We should be "evolving" into the future, through a deliberate effort to emulate the key elements of natural selection. Specifically, we need to emphasize fielding "individuals" that embody available components and "test" their fitness by performing in challenging environments. From the successes of various "experiments," we feed forward to reinforce the best components and create new ones where needed. This rate of evolutionary experimentation has to be short enough to adapt to changes in the environment and, ideally, to influence the rest of industry's own evolutionary plans. All components, from software applications and software standards, to processors, storage, and communication devices are on an approximate 12-month generation schedule these days. Nothing is permanent, only the fittest live on, and all components are of limited life-time value. By focusing on rapid *time to value*, we put the horse in front. If we instead try to nail down some particular set of "answers" to the thousands of potentially pertinent but ephemeral technology questions, we merely increase friction without gaining leverage. Even the Object Management Group now recognizes the fluidity of standards and, for this reason, has attempted to shift architects' efforts up a level, through use of MDA. In DoD, every choice needs to be open; we are running an investment portfolio, not building a cathedral for the ages. Standards and technology are the shifting sands on which we must move with agility to capture value.

The General Critique:

Few successful product lines or long-lasting complex systems have been developed without component-based, product-line architectures that generalized previous successful systems. The essence of such product-line architectures consists of an excellent understanding of key use cases or end-to-end transactions, the keen modularization of those into essential functions that combine effectively, a deep appreciation of challenging quality attributes (such as end-to-end latency), and a composition, connection and control strategy that assures that an appropriate configuration of specific implementation elements will perform acceptably in its particular target environment. In the era of web services, the same challenges remain although the infrastructure expands to include wide-area networks.

Web services make some things easier and some things more difficult, as is the history with every other approach to distributed computing. Success still depends on our being able to put together solutions from various suppliers that do important things for their important users very well. For example, we might want to accomplish time-critical targeting in less than five minutes, using various information sources, analysis and C2 tools, perhaps involving dozens of component-based capabilities and connections. Creating a system that does such important transactions well will depend upon domain-specific issues, such as: what kind of information is required; how is it represented; how should it be translated; what kinds of pedigrees are available; which analysis or C2 tools are available; what access privileges are required; what resources are available and what are the constraints; how much time can be allocated to various process steps, etc. In doing this analysis, we might discover that some of the proposed standards are helpful and that others are useless or worse. We'd probably find that more than 95% of the work would be domain-specific, and that each of the standards proposed provided little guidance or lift for the bulk of that effort.

In short, to deliver value, it's both necessary and sufficient that we demonstrate we can use networked services to do important things, using the best available off-the-shelf technology and methods. It is not necessary, nor is it sufficient, that we frame the problems or the solutions in a particular generation of technology standards coming from an immature and emerging arena such as web services. Web service technology, broadly speaking, will ultimately constitute merely one aspect of the technologies combined to solve our important problems.

Successful implementations that demonstrate information superiority, enabled by networked capabilities, are the foundations on which viable architectural approaches will be formulated. Successful, end-to-end, value-delivering processes are the "proof of the pudding." Successful network-centric applications become the grist from which valid and valuable architectural frameworks will emerge. Moreover, in immature markets, such successes strongly affect evolution in both implementations and standards. Most of these proposed standards are immature, still evolving, and untested for the kinds of resource-constrained, urgent and vital military operations that the GIG NCES should address first. The standards aren't harmful *per se*, unless one assumes they are a practical, proven framework for delivering value. They aren't. They're a snapshot in time of a rapidly evolving marketplace that's exploring new ways of interacting on the Internet.

While CTO/Software at HP, I had the opportunity to fund and oversee HP's participation in several of these standards, as well as several others that might also have

been usefully added to this list. For example, the proposed standards list omits Semantic Web standards, essential for machine-to-machine information sharing. However, all of these standards are coming in advance of important implemented solutions and represent speculative approaches to very big and diffuse problems. They should be viewed as “manual tools” in the engineers’ toolkit. When they fit an aspect of the problem well, are supported with good quality implementations, and are expected to provide stable advantages for the next few years, they probably deserve serious consideration for application. However, those qualifications significantly restrict the expected applicability of the proposed standards. Since the memorandum doesn’t include any such pragmatic qualification criteria, enforcing it uniformly can be expected to reduce productivity significantly.

Imposing too many, inappropriate, or immature standards is a bad management practice in the IT field. Standards are potential aspects of any system architecture, no more or less important than other aspects. All implementation efforts should be judged in terms of expected benefits per cost, with a significant discounting for benefits not available until some time in the future. After all, we’re in the business of spending resources to improve productivity, so we should be focusing on the biggest opportunities for delivering value and looking for ways to do that as fast as possible, while keeping costs low. Immature standards usually rate poorly, among all other considerations, in trying to implement systems today that can deliver high-value tomorrow. The most successful and productive engineering teams emphasize wise selection of and commitment to products and standards that match the implementation objectives and provide leverage for realizing the value proposition. They eschew naïve embrace of relevant but immature standards, as experience shows that most evolving software “standards” never become generally important.

One way to remedy the proposal’s greatest risks would be to consider standards compliance as one aspect of an overall project’s evaluation factors. An intelligent managerial approach would evaluate each potential project in terms of its (1) *time to value*, (2) benefit to cost, and (3) fitness for continuous evolution. Such an approach would give appropriate weight to the quality and applicability of various technologies, including standards. Because those are constantly changing, all efforts to forecast winners or to estimate long-term benefits are suspect. Uncertainty is unavoidable. Hence, **a strategy that emphasizes immediate value, affordability, and agility should outperform any alternative.** In particular, the GIG/NCES SOA proposal overemphasizes specific momentary technology standards, with the undesirable side-effect of creating a sink-hole that will absorb energy and resources through an inappropriate emphasis on and commitment to evanescent architectural precepts.

In sum, we could minimize harm through a general recommendation to adopt “appropriate standards,” leaving the specifics open and fluid. However, to maximize bang for the buck, we must shift the focus from standards to outcomes. Specifically, we should be identifying the “transactions” or “processes” needed for superior mission performance and then commit to achieving these superior results as quickly as possible, with the most appropriate and adaptable technology. We must expect fluidity and turbulence in technology, but we should measure our own efforts in terms of **value delivered**.

January 20, 2006

OMG Comment on Proposed Standards for Implementing GIG Enterprise Services

From : Dr. Richard Soley, CEO, Object Management Group

Dr. Jon Siegel, Vice President, Technology Transfer, Object Management Group

The NCES SOA WS memorandum and standards set addresses implementation of an interoperability architecture while ignoring design. Systems as large as these under consideration must be designed before they are built; otherwise it is certain that incompatibilities and inefficiencies will exist and be discovered later when repair is at best expensive and, at worst, impossible. Standards exist for design; the best of these, from the Object Management Group (OMG), feed the design semi-automatically into the downstream development workflow where it generates most or nearly all of the implementation.

For structural and behavioral modeling, the industry standard is the Unified Modeling Language (UML) 2.0. Sophisticated capabilities include nested structures, allowing a single model to represent every aspect of a system from the highest overall view to minute detail. Near the top of this hierarchy is the overall SOA-based system with its interacting major components, their interfaces, and the protocols and standards they use to communicate; zoom in on the model and the internal structure of the services emerges. Adherence to the MetaObject Facility (MOF) 2.0 standard provides model interoperability and allows models to be parsed and transformed by software; MOF compliance also enables models to be encoded into XML Metadata Interchange (XMI) for export, transmission over the network, and import into a repository or a transformation or other tool. The Model Driven Architecture (MDA), built on this foundation, makes formal the sequence of steps from platform-independent model, through choice of platform or platforms, to final implementation.

Universally recognized throughout the industry, UML diagrams are the best and most efficient way for architects and designers to communicate system design, essential for quick and error-free implementation and deployment of large systems at multiple sites spread over a large area. That these diagrams also feed into the downstream implementation workflow is a plus. The OMG also runs a program that certifies architects on their knowledge of UML.



Dr. Richard Soley

Statement of Work for Review of DISA Draft Standards for Implementing GIG Enterprise Services

DISA will provide to the W2COG Institute the draft Standards for Implementing GIG Enterprise Services for Consortium review. The W2COG Institute will then enlist leading experts knowledgeable in SOA to include architects, service providers and technology developers, and other industry standards bodies to assess and provide a critical review on how SOA standards strategy might be adopted into the fabric of NCES.

The SOA standards review will, at a minimum, consider governance issues, maturity of technology, current state of implementation and future needs, and cost, schedule and implementation risks to success. Alignment with COTS trends and investments and compliance with other SOA/GIG documents will be included."

Review will take no more than 40 days, with report delivery 15 days later.

Cost of project is 42 K. Contracting vehicle recommended is grant to W2COG Institute.

3. NORTHCOM Wireless Study Task Order

TASK ORDER THIRTY (TO30) FOR UTILITY OF COMMERCIAL WIRELESS STUDY IN SUPPORT OF UNITED STATES NORTHERN COMMAND (USNORTHCOM)

REFERENCE: The Statement of Work (SOW) for Utility Of Commercial Wireless Study In Support Of United States Northern Command (USNORTHCOM) (hereafter referred to as "the Study") is described in Section 4.0 below. The SOW is the high-level base document that defines the work to be accomplished for the Program.

1.0. PURPOSE. This document is a Task Order (TO) to the Prime Contractor outlining the necessary tasks required to deliver and support the Study. Task Orders are supplemental work orders that provide further definition of the work to be accomplished during each phase of the Program.

1.1. THE CUSTOMER. The "Customer" is USNORTHCOM SJFHQ-N located at the following address:

USNORTHCOM / SJFHQ-N
Building 2
Peterson AFB, CO 80914

1.2. THE PRIME CONTRACTOR. The "Prime Contractor" or "Prime" responsible for executing this task order is Naval Postgraduate School located at the following address:

Code 21
1 University Circle
Monterey, CA 93943-5000
Attn: Danielle Kuska
(831)656-2099
dkuska@nps.edu

2.0. TASK ORDER APPROACH. The task order approach is to conduct a four (4) month study whose task description is delineated in Section 4.0 below. This task order will require Program Management, Research, Analysis and potentially some experimentation for scalability and validation purposes.

~~The contractor will be required to deploy communications capability necessary to support deployed forces. Deployment will be directed by the customer either by air transport (C-130) or independent mobility.~~

3.0 TASK ORDER OBJECTIVE: Analyze the state of the commercial wireless networking environments to understand market trends and direction as well as current and future technology. The analysis should concentrate on capability that can be leveraged to enable fixed/mobile voice and data connectivity at the edge of the deployed network to provide

interoperability and seamless access to the USNORTHCOM collaborative information environment. Propose technology that is readily available with robust commercial base and market share; it should not rely on government support to maintain product viability in the market place. Capabilities will enable communication during quiescent daily and contingency operations with a focus towards increasing DOD's responsiveness to the Homeland Defense and Defense Support to Civil Authority (HLD/DSCA) missions.

Stakeholder users, to include military, NGO's, and first responders will be identified.

Communications for the USNORTHCOM HLD and DSCA missions fall into the following categories:

- 1) Military command and control (normally fixed garrison communications)
- 2) Interoperable mobile communications to enable seamless communications among DoD and federal mission partners
- 3) Communications support to civil authorities to maintain civilian confidence in government

4.0. TASK DESCRIPTIONS.

4.1. TASK 1. Market Research: Perform market research of commercial communications technology including wireless and interface to wired infrastructure to enable hastily formed networks. Market research should address capabilities necessary to cover quiescent USNORTHCOM, HLD and DSCA operations as well as DoD efforts to reconstitute civil communications.

4.2. TASK 2. Analysis: Analyze and assess the market research and state of commercial technology for commercial wireless and networking in reference to the ability and applicability of increasing the effectiveness of the USNORTHCOM mission. Consideration of interoperable communications, Quality of Service per Class of Service, information security and frequency and network management will be included. Analysis criteria should include ability to absorb simply new wireless communications technology over 10 year life cycles.

4.4 TASK 3. Recommendations: Provide recommendations and alternatives, to include new developments, test, and integration of current and new systems, to meet the USNORTHCOM mission. A cost benefit analysis should be included for each recommended system to allow proper ranking of alternatives. (i.e., Are the current and To Be military trunked radio solutions based on standards that permit interoperability with commercial wireless?)

4.3 TASK 4. Roadmap: Develop a 10-year roadmap that enables a capability to increase communications interoperability incorporating new commercially viable communications technology. Include analysis of the trend in USNORTHCOM mission profile as it evolves and matures. Provide Courses of Action (COAs) to implement recommendations derived in Task 3.

6.0 DELIVERABLES. The contractor will provide deliverables, status updates and briefings throughout the contract period consistent with and IAW this USNORTHCOM Statement of Work (SOW).

Deliverable	Schedule	Description
Market Research	NLT Receipt of contract + 30 days	Perform market research consistent with commercially acceptable standards. The report will be written in a mutually agreed upon customer text and briefing format.
Analysis	NLT Receipt of contract + 60 days	Analyze research and provide a written report in a mutually agreed upon customer text and briefing format.
Recommendations	NLT Receipt of contract + 90 days	A report describing the future state of communications technology with recommendations written in a mutually agreed upon customer text and briefing format.
Roadmap	NLT Receipt of contract + 120 days	A roadmap to portray courses of action to implement recommendations written in a mutually agreed upon customer text and briefing format.

7.0 FACILITY ACCESS, PERSONNEL SECURITY AND LOGISTICS. The Government will be responsible for supporting the contractor by supplying all necessary equipment, facility access, including any badges, personnel security at any work sites, and support logistics.

4. Towards a Rich Semantic Model of *Track*: Essential Foundation for Information Sharing

Towards a Rich Semantic Model of Track: Essential Foundation for Information Sharing⁷⁶

Rick Hayes-Roth⁷⁷

hayes-roth@nps.edu

Professor, Information Sciences, NPS

April 27, 2005 (v. 0.5)

Summary

Many defense, homeland security, and commercial security objectives require continuous tracking of mobile entities. The systems that perform these functions produce information products called *tracks*. A track associates observations with the mobile entity and typically includes position, velocity, and other similar attributes. One of the best current tracking systems, the Navy's Cooperative Engagement Capability (CEC) produces timely and accurate tracks for the aircraft it observes. In other domains of interest, such as seagoing surface ships, dangerous cargo and persons of interest, tracking systems are less mature and perform worse. In the near future, we will wish to share information among different tracking systems working in similar domains.

To combine information from different sources, we will need a flexible framework that can tolerate and exploit data products from different systems, although these systems employ different representations and embody different assumptions. The most basic assumptions concern what the information is intended to mean and how it is intended to be used by a recipient. *Semantics* is the term we use to refer to *meaning*, and *pragmatics* is the term we use to refer to *goal-oriented use*.

In accordance with best practices in the technology areas of the semantic web and knowledge representation, we seek to reduce the barriers to efficient sharing of information. Our approach is to identify a rich semantic model of tracks that can support multiple important functions: (1) represent a wide variety of meanings and support a broad array of pragmatic goals; (2) reduce the time and cost required to implement capabilities to reason about a new, specialized type of track; (3) simplify the understanding and importation of external sources of track information; (4) help operators describe what attributes of tracks they value in performing their tasks; (5) significantly improve our ability to combine multiple sources of track information; (6) provide a stable and evolvable base for key standards and best practices that support information sharing; and (7) improve bandwidth utilization, raising the proportion of communicated information that recipients consider significant, by delivering valued information at the right time (VIRT).⁷⁸

The proposed rich semantic *track* model will be shared widely with appropriate communities of interest. We will also recommend methods and tools for using the semantic model to address all of the objectives just discussed. By focusing on one important example of rich semantic models, we hope to provide significant near-term

⁷⁶ This work is supported in part by a contract to NPS from NAVSEA/IWS and the CEC program office.

⁷⁷ In performing this work, I have collaborated with members of the Johns Hopkins Applied Physics Laboratory who also support the CEC program office. I gratefully acknowledge their support and collaboration.

⁷⁸ F. Hayes-Roth, "Model-based communication networks for improved collaboration and decision-making," presented at 12th International Conference on Telecommunication Systems - Modeling and Analysis, NPS, Monterey, 2004. <http://doncio.ro.nps.navy.mil/icts12/>

value and also pave the way for wider recognition and adoption of this essential foundation for information sharing.

Background

Much of the evolution of information systems has focused on improving the ability of applications to share data. In recent years, the emphasis has shifted to enterprise-wide sharing of information among systems. Most recently, with the emergence of concepts such as network-centric operations, aspirations have increased⁷⁹. Now we want to be able to share information and services seamlessly across global networks of computer-based resources⁸⁰. The Internet has suggested that we should be able to draw at will from a pool of available sources and easily combine and process information as needed.

The reality of systems integration falls far short of these aspirations, however. To make systems interoperate today usually requires us both to significantly limit objectives and undertake extensive custom engineering effort. The “friction” impeding seamless interoperability arises from differences among the participating systems, including the types in Table 1.

As Table 1 shows, two systems can differ in many ways. Most differences arise because the system developers made many assumptions appropriate to the original context for operating their particular system. Usually these assumptions were implicit, and often they aren’t even documented. Many correspond to developers’ “common sense” or conventions of contemporary engineering practice. In the U.S., many systems use dollars for currency and British units of measurement. Often these units are not explicitly represented. In the E.C., countries used to employ national currencies in addition to the metric system. Today, most European countries are adopting the Euro for currency. Occasionally, such system incompatibilities produce disasters, such as a mission to Mars that is destroyed because of measurement system incompatibilities. More often, the costs of incompatibilities are buried in the behind-schedule, over-budget integration projects that occur time after time. An entire industry has arisen in the commercial arena to address such Enterprise Application Integration engineering jobs, and this is a costly, labor-intensive business.

In national defense and homeland security, the challenges are every bit as great and the industrial practices no better. “Best of breed” systems are those that do specialized functions better than all alternative products. While we would like to combine these easily into overall, unsurpassed “systems of systems,” this proves very costly. These best of breed specialists are never designed, from the outset, to work compatibly with every other potential federation partner. When called upon to make two systems interoperate in ways that had not been anticipated originally, the engineers go through the categories of differences and apply as many mitigating methods as required to bring the system of systems up to an acceptable level of performance. Systems of this sort always suffer, however, from an increase in overall uncertainty and error, because we lack powerful methods to assure that the semantics and pragmatics of the integrated system recognize and correctly handle all important situations. In all such integrated systems, we ultimately

⁷⁹ Dave Alberts (*et al.*), OSD/ASD(C3I), CCRP, “Network Centric Warfare: Developing and Leveraging Information Superiority” (2nd ed.). http://www.defenselink.mil/nii/NCW/new_0801.pdf.

⁸⁰ Department of Defense, Directive 8320.2 (December 2004), ASD(NII)/DoD CIO, “Data Sharing in a Net-Centric Department of Defense.”

rely upon trial-and-error discovery to reveal important problems and then address them manually, one by one.

Table 2. Sources of friction that make interoperability difficult.

Type of Difference between Systems	Resulting Difficulty	Typical Method of Mitigating Problem
Data representation	Identical bits differ in meaning	Translate to a standard representation
Data precision	Incompatible approximations	Reduce to minimum precision
Data measurement systems	Incompatible units and coordinates	Translate to a standard reference
Temporal calibration	Presumed simultaneous data originate at different times	Combine information with more uncertainty
Geospatial calibration	Presumed identical positions vary in space	Combine information with more uncertainty
Attributes and scales	Similar but non-identical aspects assessed differently	Translate to a common framework and heuristically combine
Concepts	Similar names used for non-identical classes and relations	Combine information with more uncertainty
Events & Triggers	Similar names used for non-identical conditions	Combine information with more uncertainty
Processes	Similar names used for differing states, conditions, and actions	Translate to a common framework with more uncertainty
Resources	Similar names for differing resources, costs and policies	Translate to a common framework with more uncertainty
Evidence, Association, Inference, Belief and Uncertainty	Different inductive methods to relate and combine evidence and to support inferred beliefs	Adopt one preferred approach, adapt the compatible information to it, and drop the incompatible information

We have basically three approaches to integrating systems and sharing information today: the standardized data-centric approach, the human translator approach, and the common intermediate hub-and-spoke approach. The first method standardizes all aspects of data, across all functions and applications. A single unified data model is created for all purposes, and all applications use it consistently. The second method relies upon people to translate information from one system into another one, thereby relying on human understanding to make appropriate mappings between underlying semantic models and associated pragmatics. The third approach is to create new standard information models that become the hub of a hub-and-spoke like interchange. Each system that produces information can publish and share its data using the same hub model for all recipients. In

this approach, each system translates its understanding of its own information into the semantic and pragmatic framework of the hub model. Likewise, each consumer of information finds relevant information through the hub and translates it into its own representations consistent with its understanding of the intended semantics and pragmatics.⁸¹

Each of these approaches has significant deficiencies though. The first method is slow and brittle, because it requires every system to accord with a single integrated semantic and pragmatic model. Creating such a model can take forever, and it cannot evolve as rapidly as needs for new capabilities arise.

The second method is labor-intensive, knowledge-dependent, slow and error-prone. In this case, we are asking people to do routine, repetitive data reading and writing tasks, often between systems whose semantic and pragmatic assumptions they don't know well.

The last method allows systems to develop in parallel, but it presupposes that the hub provides a single integrated semantic and pragmatic model. If multiple models exist, then every publisher needs to have its information specially translated to meet the contextual requirements of every consumer. The publisher doesn't know all of the consumers that well. The consumers don't know all of the publishers that well either. In any case, new publishers and new consumers continually enter the arena, and there's no way for the expanding set of required translations to be melded into a single and stable hub.

It's no wonder then that most efforts to create seamless systems of systems remain pipedreams today. Our approaches to these lofty goals are reminiscent of people who want so much to reach the sky that they climb up the tallest tree as fast as they can. Instead of working harder, we need a radically different approach that can meet the true challenges of sharing information among systems that embody different semantic models because of different pragmatic concerns and operating contexts. Table 2 below identifies the principal requirements for efficiently sharing information among such systems.

As Table 2 shows, we desire numerous qualities of the systems that share information. We want them to understand the meaning of data that arises from different contexts with different pragmatic concerns. We want to combine information in sensible ways. We want our systems to improve continually, because it's impossible for them to be born perfect. Furthermore, since the sources and purposes change continually, we want our systems to be able to exploit new sources automatically, adapting to their associated semantics as appropriate. In the middle column of Table 2 we identify the principal functional capabilities that would achieve these desired qualities. Then, on the right-hand-side of the table, we list the technical strategy proposed to implement each of these capabilities.

⁸¹ In each of these approaches, many important details are being glossed over. In the hub-and-spoke model, for example, the hub may be actual or virtual. Particular publish-subscribe systems may take different forms while accomplishing equivalent results. In many cases, it will make sense to have more than one hub, supporting each active community of sharers. Lastly, it might prove advantageous in some settings to create a multilingual translator that specializes in the hub's semantic model and that can provide translation services for many different suppliers or consumers of information between their particular languages and the hub..

Table 3. Principal requirements for information sharing.

Desired Quality of Information Sharing	Capability Required to Achieve Desired Quality	Basic Strategy to Achieve Required Capability
Employ different semantic models	Read and interpret a semantic model	Use models based on formal meta-models with grammars
Understand the meaning of data	Parse data into its associated semantic model	Use models based on formal meta-models with grammars
Express meaning as data correctly	Ability to generate data from its associated semantic model	Use models based on formal meta-models with grammars
Translate meanings between two systems	Map from one semantic model to another	Use models based on formal meta-models with grammars
Understand what information tasks need	Read and interpret a model of task pragmatics	Use models based on formal meta-models with grammars
Verify that required information is expressible	Assure a semantic model supports the required pragmatics, or continually improve it	Map required conditions for actions into corresponding semantic expressions
Tolerate and exploit diverse sources	Operate simultaneously with multiple models	Maintain segregated namespaces as required
Handle ambiguity intelligently	Recognize ambiguities, adapt to them, and continually improve	Detect and reduce 1-to-many translations
Handle inconsistency intelligently	Recognize inconsistencies, adapt to them, and continually improve	Detect and reduce logical impossibilities
Handle errors intelligently	Recognize errors, adapt to them, and continually improve	Accept negative feedback, trace and reduce causes
Handle quality variations intelligently	Recognize differences in source qualities, adapt to them, and continually improve	Use inconsistencies and errors to reduce reputation of responsible source
Combine information appropriately	Collate correlated information into coherent association sets	Bundle assertions that are and are not consistent
Exploit and eliminate redundancy	Employ heuristic methods to reduce correlated information	Implement best methods of empirical inference
Justify results derived from various sources	Track information pedigree, as required	Retain histories of inferences and sources

The key strategic ideas in Table 2 are summarized here:

- Use models based on formal meta-models with grammars, and automatically generate required input and output language systems.
- Map between models as needed, especially when assuring that the semantics are adequate to support the pragmatics.⁸²
- Combine, reduce, and track⁸³ information as appropriate.
- Continually improve by recognizing problems and changing knowledge to reduce or eliminate them.

The purpose of this paper is to introduce the problem and the strategic approach rather than providing a long, technically detailed treatment. Such an extensive technical treatment is premature, anyhow. We don't currently possess well defined tools and methods for this work. That is where we want to go.

In order to make these abstract and ambitious approaches more understandable, in the next section we'll delve into some examples of important goals we have for tracks and will draw upon one of the best defined semantic models of defense concepts, namely those included in the C2IEDM model created by NATO's Multilateral Interoperability Programme (MIP)⁸⁴.

Semantics and Pragmatics of Track, by Example

Tracks are an important element of situation assessment in most command and control systems. In ground combat, commanders need to determine where enemy forces are, how to avoid them, how to counter their attacks, or how to attack them while they're stationary. In air combat, similar decisions must be made and corresponding actions taken. Ground vehicles move at speeds between 0 and 100 mph. Air vehicles moves at speeds up to Mach 3 or so, although most move at speeds between 60 knots and 600 knots. Surface ships move at speeds normally under 40 knots, though some small ones can go faster. Dismounted infantry moves at speeds under 10 mph. In all cases, commanders want to track these, anticipate their likely motions and potential threats, determine how best to counter threats, and then implement chosen countermeasures efficiently.

From these specific examples of differing mobile entities and general pragmatic concerns, we can identify the following common pragmatic objectives for a mobile entity M with possible intentions and capabilities to do harm to our interests:

1. Observe, detect, identify, classify and continuously monitor M.
2. Infer M's intent.

⁸² W. Ross Ashby coined the famous "law of requisite variety," which basically states that any system must perceive situational distinctions sufficient to enable it to make appropriate differentiated responses required for success in its environment. Our requirement for systems that combine information is similar: the semantics must be sufficient for the pragmatics. Ashby, W. R. (1958), "Requisite variety and its implications for the control of complex systems." *Cybernetica*, 1(2), pp. 83-99.

⁸³ The word "track" in the above bullet is used to mean "keep a record of its origins and the processes that converted inputs into new products." Such a record is sometimes called a "trace" or a "pedigree." The rich semantic *track* model that's the focus of this paper uses the defense domain concept of a "track" to mean the product of observing a mobile entity to identify it, monitor it, and predict its behavior. We will italicize this meaning of *track* throughout the paper.

⁸⁴ See <http://www.mip-site.org/>.

3. Determine M's threats $T_{M,D}$ against domain D.
4. Locate M.
5. Predict M's location and behavior.
6. Alert agent A about M and threats $T_{M,D}$.
7. Determine countermeasures $CM(T_{M,D})$ to threats $T_{M,D}$.
8. Inform agent A about countermeasures $CM(T_{M,D})$.

These eight pragmatic objectives define the general and common concerns of military and security agencies with potentially dangerous mobile entities. The whole purpose of sharing information among different sources is to support these common objectives. The premise of this paper is that we can best achieve that purpose by relating all information sources to those purposes. While individual processes might differ for each of these concerns, we should be able to express what the information requirements are for each process in terms of semantic capabilities. Further, we should be able to create translators to re-express various sources of information in terms of a hub semantic model that provides the capabilities our pragmatic processes require.

Let's consider how this can be done. To do this, we will write pseudo-code in the style of Prolog rules. Each rule will be of the form $C(x, \dots, z) \leftarrow P(x, \dots, z) \& \dots \& R(x, \dots, z)$, with the following interpretation: To infer or conclude that $C(x, \dots, z)$ is true, it suffices to conclude that $P(x, \dots, z)$, \dots , and $R(x, \dots, z)$ are all true. The variables x, \dots, z may be replaced by any specific term, as long as substitutions are done correctly.

As a simple illustration, we might have a rule that says that a commercial aircraft from one's own country, observed to be following its planned route, has the intention of completing its filed flight plan. We might write this roughly as follows:

[Rule I1] $\text{Intention}(M, \text{Follow-its-filed-flight-plan}, \text{High-confidence}) \leftarrow$
 $\text{Commercial-aircraft}(M) \& \text{Affiliation}(M, \text{"U.S."}) \& \text{Following-planned-route}(M,$
 $R) \& \text{Current-planned-route}(M, R)$

Conversely, we might assume a hijacked aircraft has a variety of possible intentions, including using the aircraft as a missile to attack some target or diverting to a location not on the original planned route and landing there. Such an inference might be written in terms of two rules such as these:

[Rule I2] $\text{Intention}(M, \text{Fly-into-a-target}(M, t), \text{Probable}) \leftarrow$ Commercial-
 $\text{aircraft}(M) \& \text{Hijacked}(M) \& \text{Target}(t) \& \text{Can-reach}(M, t)$

[Rule I3] $\text{Intention}(M, \text{Deviate-and-land}(M, da), \text{Probable}) \leftarrow$ Commercial-
 $\text{aircraft}(M) \& \text{Hijacked}(M) \& \text{Airport}(da) \& \text{Can-reach}(M, da)$

Rule I2 states that M intends to fly into an unspecified target t that it can reach. Similarly rule I3 states that M intends to deviate to and land at an airport da, where da is undetermined. The airport da is one that M can reach with available fuel. Both rules I2 and I3 state that the inferred intentions are "Probable," in contrast to rule I1 which rates its inferred intention as High-confidence.

Any formal system for expressing rules such as these must follow some syntactic conventions. Here we've used the convention that lowercase terms are unbound variables that can ultimately be instantiated by specific constants. Uppercase terms, on the other hand, are constants that name various entities or concepts. For example, the constant

Probable stands for a degree of confidence or belief that is judged more likely than either impossible or unlikely.

Any system of concepts will have its own nuances and best practices for modeling the world effectively. Our assumptions are that no system is perfect, perfection is both an unachievable and unwise goal, and that great benefits can derive from creating workable systems that significantly improve our speed and effectiveness. Therefore, while we could dwell on different approaches to representing each concept and reasoning about it logically or empirically, we won't do that here. Instead, we wish to initiate use of evolvable semantics to support important pragmatics. Thus, the key capability we need is to do some things well while being able to improve continually. For that reason, almost any reasonable semantic system will be good enough for significant information sharing. The essential quality required is that the system distinguishes states that warrant different inferences and actions. In the above rules, for example, the predicate Hijacked(x) distinguishes a state sufficient to support different inferences about the intentions of the aircraft. Any system that makes distinctions that correspond to Hijacked(x) can be used through translation for the same pragmatic purposes.

What the examples show is that pragmatics aims at performing important functions, such as the eight general ones cited above. Each of these objectives requires inference and problem-solving to assess available information and determine which inferences are warranted. The information required, initially, is *conceptual*, rather than particular or concrete. For example, one type of information required was "is an aircraft hijacked?" This is a question about the state of the world or, more precisely, about beliefs about the true state of affairs. Different information systems will represent and store such beliefs in different ways. What is necessary is that available information pertinent to this conceptual requirement is mapped it, somehow, so that the inference process can proceed as appropriate.

Given a set of pragmatic objectives, the inference process relies upon conceptual categories. A semantic hub should make all of the conceptual distinctions required to support those categories and related pragmatics. The rich semantic *Track* model, therefore, should reflect aspects of state that most users of track information require for addressing expected pragmatic concerns. As we employ such a model to intermediate sharing among systems, we will inevitably discover additional concerns not yet adequately addressed in the current model. This will drive an iterative, evolutionary series of improvements to our model of *Track*.

Table 3 below enumerates many of the required concepts to support the eight principal types of general-purpose pragmatics for *Track*.

The most important point from Table 3 is that pragmatic concerns regarding *Track* are fairly generic, stable, and procedural. We should be able to create a mostly-hierarchical conceptual scheme working backwards from pragmatic objectives to required concepts to supporting distinguished data values. The ability to adapt this standard hierarchy rapidly to exploit a new source would be the operational test of value. This suggests both what types of products we need and also what types of methods will enable us to adapt these products to new situations. In the next section, we provide a sketch of the semantics that should provide the required scaffolding for this approach.

Table 4. Semantic concepts required to support *Track* pragmatics.

Pragmatic Goals	Assumptions and Inferences	Required Concepts	Semantic Aspects
1. Observe, detect, identify, classify and continuously monitor M.	Continuity of motion; Physical persistence; Observability; Immutable Identity	Geospatial and temporal coordinate systems; position; velocity; behavior history; classes, types and identification	Position and dynamics in relation to reference system; measurement systems, registration, and errors; behavior, states, and state transitions
2. Infer M's intent.	Mobile platforms are controlled by pilots; pilots normally follow filed plans; hijackers take control of hijacked vehicles; hijacking blocks normal behavior; hijackers have abnormal intentions	Plans; filed plans; modified plans; persons in control; operators of vehicles; hijacking; hijacker; abnormal intentions; expected behavior; discrepancy from expected behavior	Plans as intended states and behaviors with dependencies and constraints; actors and resources in plans; actors' goals, intentions, plans and resources; planned effects as expectations; behaviors consistent with and inconsistent with expectations
3. Determine M's threats $T_{M,D}$ against domain D.	Domains include property, states, cultures, people; threats to domains are possible ways to do harm to elements of the domain; the more probable and hurtful the damage, the worse the threat	Domains and their resources, symbols, systems, centers of gravity and key attributes; types of harm; harmful processes; potential harm; probable harm; expected harm; capability to inflict harm; threat	Important entities and important attributes of them from a security perspective; vulnerabilities; attack methods and profiles; estimated success and consequences of attacks; time and other remaining barriers to the success of the attack
4. Locate M.	Ability to change location limited by maximum velocity and physical constraints; objects in motion most likely to continue consistent with historical behavior; identity required for location	Reported position at time t; Inferred position at time t; Probability density around position or other confidence intervals; confidence about identity	Position and dynamics in relation to reference system; estimates of error and uncertainty about position, dynamics, and identity

5. Predict M's location and behavior.	Use same inferences in 1 adjusted to reflect expected behaviors (from 2, 3 and 4)	Current state relative to executing plan; expected behaviors and variations	Variations in speed and rate of climb based on phase of flight; probable variations in execution;
6. Alert agent A about M and threats $T_{M,D}$.	Some agent is responsible for knowing about vehicles including M or threats including those in $T_{M,D}$; each agent has preferred ways of being alerted	Areas of responsibility; agent identities; means of communicating; agent sensitivities and preferences; agent context, state and focus; required quality attributes of reports	Roles for monitoring routine behaviors, abnormal behaviors, and threat behaviors; assignment of agents o roles; communication methods and required protocols; current beliefs of agents; value of various belief changes
7. Determine countermeasures $CM(T_{M,D})$ to threats $T_{M,D}$.	For known threats with predetermined countermeasures, collect those and instantiate them; for others, determine ways to block attacks	Threat types and countermeasure types; variables and substitution; ways to block causal chains by denying essential prerequisites, etc.	Destroy an enemy vehicle prior to its attack; impede passage through a requisite space; eliminate resources required to sustain the attack; covertly move the target or mask with a decoy
8. Inform agent A about countermeasures $CM(T_{M,D})$	Use methods in 6, appropriately adjusted to address the agent responsible for countermeasures.	As in 6, adjusted appropriately.	As in 6, adjusted appropriately.

An Initial Specification of Track Semantics

The first step in developing a rich semantic model of *Track* is to determine how it can support the pragmatic requirements, as indicated in Table 3. From that analysis, we can see that the *track* model must allow us to describe our beliefs about a mobile entity and its past, present and predicted future states. In addition, we will need to be able to justify inferences that we make as part of the tracking process. So the track model will necessarily consist of two principal types of information, one that describes our beliefs about the tracked entity and another that describes the qualities of those beliefs. This is an example of information and meta-information.

Before giving a formal specification for this belief structure, it makes sense to present the structure as a conceptual hierarchy, introducing the names we propose for different categories of information and the relationships among these categories. In a conceptual hierarchy, which is much like a topic outline, the most general concepts are the outermost items of the outline. Successively indented topics represent specializations or subcategories under the topic they descend from. To illustrate these points, consider the following abbreviated hierarchy shown in Figure 1:

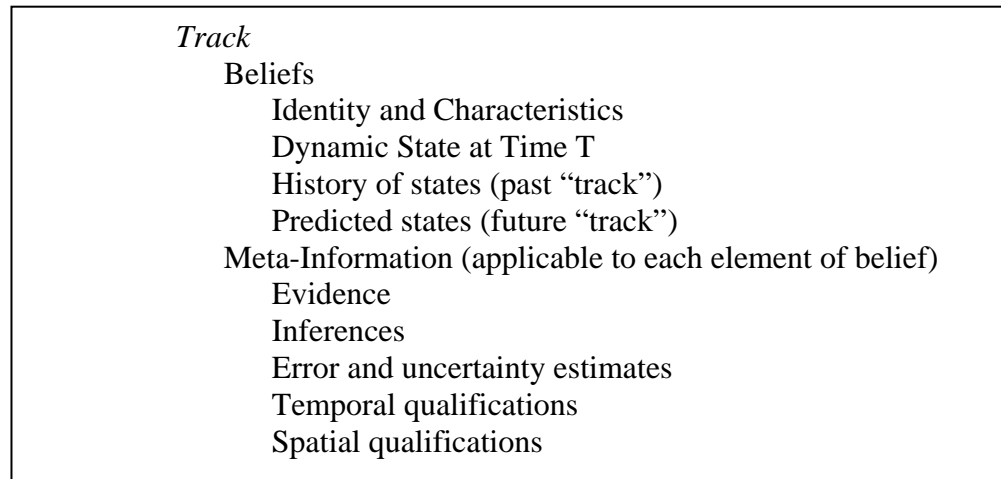


Figure 6. The top-level conceptual hierarchy for *Track*.

This fragment of a conceptual hierarchy describes the most general, or topmost, element called *Track*. The concept *Track* contains two principal component concepts, called *Beliefs* and *Meta-Information*, respectively. The parenthetical note on *Meta-Information* indicates that all of the components of *Meta-Information* may apply to each element of *Beliefs*. This indicates that when we use the conceptual hierarchy to create actual beliefs that are instances of *Track* *Beliefs*, we may find it useful to qualify every belief by using the sub-concepts of *Meta-Information*. While this creates an opportunity for significantly expanding the volume of data in our systems compared to traditional databases, the purpose of the semantic model is to make possible precise description of potentially important states. In general, our systems will create sparsely instantiated models, because many aspects will not be deemed relevant or material. Furthermore, most implementations will reduce the total volume of data by finding ways to compress and abbreviate bodies of information with inherent redundancies. In short, we should focus now on *what* needs to be represented rather than on *how* data can be stored and accessed efficiently.

The first-level sub-concepts of *Track* are Beliefs and Meta-Information. By convention, I will use the plural form of English nouns or, in the case of Meta-Information, a collective noun to indicate that there may be any number of instances of those concepts in an actual application. Thus, we should expect in any given track, to find one or more beliefs. These beliefs will, in turn, instantiate the sub-concepts of Beliefs, which means they will describe the Identity and Characteristics, the Dynamic State, the History of states, or the Predicted States of the tracked entity. Typically, we should anticipate that each tracked entity will be described by one belief of each of these four types when the entity has been well identified and confidently tracked. For example, in a civilian air traffic control scenario, a general aviation aircraft following an IFR (instrument flight rules) flight plan would be so described. Specifically, the Identity and Characteristics would be the aircraft's registration ("tail number"), type, make and model, and navigation capability (such as GPS-enabled). Its Dynamic State at the current time would describe its coordinates, altitude, heading, airspeed, groundspeed, number of passengers, fuel remaining, next waypoint, destination, and assigned transponder code, among other dynamic features. The History would contain a past record of such dynamic states, thereby enabling us to review where the plane had been and how it had traveled to its present state. The Predicted states would, for example, indicate expected arrival times at each of the waypoints along the planned route of flight. These predicted arrival times presumably would be updated to take into account effects of winds on groundspeed as well as anticipated changes in groundspeed that climbs or descents in future route legs will cause. In this way, people tracking the flight will have expectations, with some margins of uncertainty, around times that the plane will cross specific points along the planned route. Additional predictions could reasonably be made by interpolating between these specific predicted states.

The simple example above was intended to illustrate how the top-level elements of Beliefs should be used to describe a simple tracked entity. One of the objectives in formulating this rich semantic model of *Track* is to make it easy for developers to create compatible tools and for users to describe the important aspects of the tracks they are concerned with in their particular domains of application. The well-identified, accurately tracked, pre-planned general aviation aircraft is probably the simplest possible case. As we increase the complexity of the entity, the uncertainty about its intentions, the errors in observation, and the challenge of predicting its future states, the information we must create becomes more complicated, uncertain, and voluminous. The same *Track* model should be able to represent all important aspects of beliefs in these cases as well.

The second major sub-concept of *Track* in Figure 1 is Meta-Information, and this includes Evidence, Inferences, Error and uncertainty estimates, Temporal qualifications, and Spatial qualifications. The meaning of each of these categories is as follows. Evidence describes the observations, sensors and reporters that back up the belief. Inferences describe the additional beliefs that should be inferred from the initial belief, because they follow logically or empirically. Error and uncertainty estimates specify limitations and qualifications on the belief that should constrain how we employ the belief. Temporal qualifications further constrain our use of the belief to particular intervals of time when the belief is expected to be valid. In a similar way, use of beliefs might be constrained to particular regions of space, and Spatial qualifications specify such constraints. For example, a forecast for thunderstorms typically applies to a particular region. As another example, we might want to limit the predicted location of an aircraft to a range of altitudes between 200' above ground level and 18,000' if the aircraft is flying VFR (visual flight rules).

Now just to review briefly the main point of this paper, we want to use a semantic model of *Track* to make it easy for systems to share information. Using just the high-level concepts of Figure 1, this suggests that consumers of track data typically want to know one of four types of beliefs: (1) identity/characteristics; (2) current state; (3) past state; and (4) future state. Furthermore, since all beliefs are necessarily limited by the quality of the evidence and other limitations on when or where they can be appropriately employed, consumers want to know what qualifications/limitations apply to whatever beliefs suppliers provide to them. The point of this paper is that we can accelerate and improve sharing by providing a common basis for expressing these concerns so that all suppliers and consumers of track information can rely upon it as a semantic hub used in interpreting information from different systems⁸⁵. Of course, the same approach should be useful for other types of information than *Track*. In the end, the *communities of interest* (COIs) that are concerned with each important concept should take ownership of the process by which key concepts are formalized, how translation is operationalized, and how inevitable evolutionary changes are supported. Our focus on *Track* should provide a solid foundation for an important class of sharing. Many elements of the semantic model should generalize to other concepts and other domains. In particular, the basic structure of Beliefs and Meta-Information should prove widely applicable and robust.

Implicit in this general strategy for information sharing is the idea that members of COIs should find it *easy* to read and understand instances of the semantic models relevant to them. While simply stated, **this is a profound objective**. First, if practitioners find it easy to read and understand semantically rich information, this will reflect that the information has been structured and presented to them in ways they find natural and efficient. Experts and skilled personnel in nearly every domain of interest have developed ways of structuring and presenting information that simplify and improve their performance. Tracks, for example, are often shown as current location points with predicted future vectors. They may also include past positions as a series of connected points. Each track might be color-coded in a distinct way to indicate other key characteristics or simply to reduce confusion. All of these lines are superposed on a map, of a selected space and time, to make it easy for the operator to grasp quickly the state of affairs. In air traffic control, for example, the controller transfers this information into his own mind where he or she mentally maintains the dynamic model, often called “the bubble.”

So there are two aspects of the profundity alluded to above. First, regardless of the detailed, perhaps complex, semantic model that underlies the information being maintained in our systems, the human’s *view* is often highly tailored. The view converts information from representations that designed for use by computerized reasoners to forms that are quickly and effectively grasped by humans. The second profound aspect is that the usefulness of a semantic hub isn’t fully evident until these human-oriented viewers and editors are in the hands of operators who demonstrate their value. Semantic models should make it possible to create better viewers and editors more easily, and should also support the need to continually evolve and improve them. Often in the history

⁸⁵ While it would be nice if only one hub were required for all purposes, this seems unlikely. A semantic hub is, in essence, a “language” to support a community of people who share information because of overlapping interests. It’s rare that one language supports all interested parties. There are many reasons why we routinely find multiple languages among peoples, even when they have common interests. Two concepts are fundamental in the anticipated process: (1) any community of interest should be able to have its own language and should be able to control its evolution; (2) communities with overlapping interests will constitute a higher-level aggregate community that, in turn, will need to use an agreed semantic hub to interoperate. Thus, languages will evolve, within and between communities, probably forever.

of information systems, the first excellent viewers and editors come into existence as part of a proprietary, integrated stovepipe system. Later, as technology evolves, data is separated from code, and eventually data is more explicitly modeled semantically. This allows competitive approaches to be pursued for viewing and editing information.

Semantic hubs for important information will accelerate the development of superior viewers and editors for humans. This will pay double dividends. In addition to improving their ability to grasp and exploit information productively, it will also enable practitioners to identify quickly and effectively shortcomings in the information itself, hence accelerating the debugging and evolution of the semantic models. As an important consequence, the rate of continuous improvement in information sharing will increase.

The full proposed *Track* conceptual hierarchy is described in Table 4. In this table, the level of indentation indicates degree of subordination in the hierarchy. That number is shown explicitly in the first column. The second column contains the corresponding concept. The last column briefly specifies the meaning of the concept.

Table 5. Concept hierarchy in *Track* semantic model.

Level	Concept	Notes on Meaning
1	Beliefs	Collection of believed assertions
2	Identity and Characteristics	Identifies the tracked entity
3	Owner	Who owns the entity?
4	Affiliation	What is the owner's affiliation?
3	Operator	Who operates the entity?
4	Affiliation	What is the owner's affiliation?
3	Registration number	The entity's registration number
3	Communications call sign	The entity's call sign
3	Weight	The weight of the entity
3	Observable features	Features one can observe
3	Aggregation	Is the entity a set of entities?
4	Components	What are the contained entities?
4	Structure	How are they structured?
3	Construction features	How constructed and of what?
3	Class	Broadest classes of vehicles
3	Category	Broad sub-classes of vehicles
3	Type	Specific make & model of vehicle
3	Capacities	How much can it carry?
4	Fuel	How much fuel?
4	Load	How much weight in total?
3	Capabilities standard for type	Specifications for the vehicle
4	Take-off	Parameters about its take-off
4	Landing	Parameters about its landing
4	Range	Parameters about max.trips
4	Altitude	Parameters about its altitudes
4	Speed	Parameters about its speeds
4	Operation in icing conditions	De-icing capabilities?
4	Maneuver	Turns, maximum loads, etc.
4	Evasion	Capabilities for evasive action

4	Stealth	Capabilities for stealth
4	Defense	Capabilities for self-defense
4	Offense	Capabilities for inflicting harm
4	Support	Capabilities for supporting others
4	Diversion	Capabilities to divert from plans
4	Turnaround	What's needed to start new plan?
4	Differences from standard type	Special attributes of this entity
3	Operational Characteristics	How it performs normally
4	Limitations	Constraints on operations
4	Resource requirements and consumption	Resources consumed for various activities
2	Dynamic State at Time T	Values of variables, at time T
3	Time T	The time index T of the state
3	Temporal coordinate system	The coordinate system for T
3	Position, velocity, acceleration, etc.	Location and time rate of changes
4	Spatial Coordinates	Fixes on various dimensions
5	Spatial Coordinate system	Dimensions and origin
4	Error of measurement	Estimate of uncertainty or error
3	Operations	How it's being operated
4	Control	Who or what is controlling it?
5	Possessor	Who or what is in possession?
6	Affiliation	What is his/her/its affiliation?
5	Crew	Who are the crew members?
5	Non-crew	Who else is augmenting the crew?
4	Intent	What is the controllers' intent?
5	Peaceful	Is it peaceful?
5	Threatening	Is it threatening? How?
4	Plan	The plan to achieve intent
5	Route	The route the plan incorporates
6	Waypoints	Key points along the rout
6	Corridors	Key pathways between waypoints
5	Timing	When do key events occur?
5	Tactics	What tactics will be used?
5	Resources	What resources will be used?
6	Personnel	What personnel will be used?
6	Consumables	Other resources consumed
6	Systems	Other systems required
6	Weapons	Weapons required
4	Carried load	Load that is actually carried
5	Electronic equipment	Electronic equipment carried
5	Weapons	Weapons carried
5	Crew	Crew members carried
5	Passengers	Passengers carried
5	Cargo	Other cargo carried

4	Other Behaviors	Other behaviors that can occur
4	Other Qualities	Other qualities that characterize or limit operations
4	Transponder code	The identifying code assigned
4	Dynamic variations in Identity and Characteristics (if any)	Changes over time in typically static attributes
3	Owner	A changed value for owner
4	Affiliation	A changed value for affiliation
3	Operator	A changed value for operator
4	Affiliation	A changed value for affiliation
3	Registration number	A changed value for registration
3	Communications call sign	A changed value for call sign
3	Weight	A changed value for weight
3	Observable features	A changed set of observables
3	Aggregation	A change in overall composition
4	Components	A change in units aggregated
4	Structure	A change in the structure or organization of aggregate
3	Construction features	Changed materials or manner
3	Class	A change in the broadest class
3	Category	A change in the broad category
3	Type	Changed value of defining aspect
3	Capacities	A change in its capacities
4	Fuel	A different amount of fuel
4	Load	Changed value of maximum load
3	Capabilities non-standard for type	Type properties changed
4	Differences from standard type	Ways different from type changed
3	Operational Characteristics	Changed operational properties
2	History of states (past “track”)	Past dynamic states for entity
2	Predicted states (future “track”)	Future dynamic states for entity
1	Meta-Information (applicable to each element of belief)	Information about and in support of the referenced belief
2	Evidence	Evidence supporting the belief
3	Observations	Data from supporting observers
4	Reported values	Actual reported values used
5	Time of observation	When observation occurred
5	Sensor	Sensor making the observation
6	Capabilities	Capabilities of the sensor
6	Dynamic state at time of observation	State of dynamic aspects of the sensor when observation made
4	Reporter	Who or what reported the data?
3	Evidentiary events	Evidence from related events
4	Confirming events	Events that confirm the belief
5	Confirming events to notice	Events predictable from belief
5	Confirming events detected	Predicted events observed

5	Confirming events missed	Predicted events not detected
4	Disconfirming events	Events incompatible with belief
5	Disconfirming events to notice	Events predicted to not occur
5	Disconfirming events detected	Events observed contrary to belief
5	Disconfirming events missed	Events <i>not</i> observed, as expected
3	Related beliefs	Other beliefs related to this one
4	Incompatible beliefs	Beliefs not mutually possible
4	Implied beliefs	Beliefs implied by this one
4	Implying beliefs	Beliefs that imply this one
2	Inferences	Beliefs inferred from others
3	Quality of inferences	Justification and assessment
2	Error and uncertainty estimates	Estimates of error and uncertainty
2	Temporal qualifications	Limitations on time of belief
2	Spatial qualifications	Limitations of location of belief

Table 4 provides a table of contents for the *Track* model. It leaves out many details, such as the specification of permissible values and constraints among them. It also doesn't say which of many alternative systems of measurement and description would be best for any of the various values. For example, there are several different standard frameworks for geospatial measurement and location. While it may prove useful to select one best coordinate system for the first semantic hub, our techniques should be open to the use of multiple alternative systems for any of the conceptual elements. It's not really important which system is used, so long as we support translations into and out of the hub to meet the requirements for getting suppliers' information to consumers in the form they need it.

Furthermore, since the most constrained resource is likely to be the skilled human operator, the key factors in determining the effectiveness of any information sharing environment will be the naturalness and ease of use of the viewers and editors provided to the operator for each kind of information. Operators most familiar with one type of coordinate system should be able to view information from that point of view. There may be some loss of precision when translating between different frames of reference, but this increase in error or uncertainty should be easy to depict and explain. The most important thing is to relate all information to some hub model or models that can enable consumers to get the sources translated to the form they find most productive to work with.

Using the Track Model to Achieve the Stated Objectives

We will leave the details of how best to represent legal values and constraints among them for a future paper. There should be no doubt that all of the types of concerns included in the *Track* model can be expressed in terms of some grammar of permissible constructions and that legal values and other constraints can be expressed similarly. This means that we will be able to have at least one hub semantic model for the entire *Track* concept. If we find it desirable to allow multiple alternative formalisms within the hub, that also presents no special problems. Each component of the hub model must, however, be supported with two types of tools so the overall approach can prove valuable: (1) viewer/editors are required so people can create, understand and modify *Track* Beliefs

and Meta-Information; (2) translators must be written to and from the hub model to support valued consumers and suppliers of information.

Once these two types of tools are available, the objectives are at hand. First, we make it possible for operators to access information supplied by others. The basic method used is to allow operators to specify the type of information they seek and then to provide relevant information to them. To specify what they seek directly, the operators can formulate “queries” using a viewer/editor adapted to the purpose of expressing information “goals.” A conventional approach provides forms that allow the user to fill in values of sought variables and other unconstrained variables that would be associated to these. Information is sought which matches the constrained variables, and the values of all the associated variables are presented to the user. In database systems this is often called Query By Example (QBE). Other query languages are also readily applicable.

The user’s queries, stated in terms of the user’s view, must be translated into equivalent queries expressed in terms of the hub semantic model. That is done by one of the translators. A query planner then determines which aspects of the query to dispatch to various available information systems, based on its knowledge about the efficiency of asking various queries to the various systems. The query planner then would translate each sub-query it intends for each supplier system into a corresponding query expressed in terms of the semantic model of that system. One of the translators does that, translating hub semantics into the specific semantics of the targeted system. Once query responses are produced from each system, they need to be translated back into the hub semantics, combined by the query planner into an answer, and then translated back into the semantics of the operator’s chosen environment, where the operator’s preferred viewer/editor would display the requested information.

The methods of the preceding paragraph address how answers to queries are found in existing information stores. A variant of this approach must be used to create monitors that will alert users when new information becomes available. This may mean setting up “standing queries” for database systems or selective filters for publish-and-subscribe systems. The essential operations are similar in this case. The standing queries or filters are best expressed in terms directly supported by the information systems used by the suppliers. In such a case, the queries are first formulated in user view terms, then in hub semantic terms, and finally in supplier system terms. Answers to queries are translated back through the reverse steps. Of course, different system environments might lead us to do translations and query processing in different orders, but these variations seem straightforward.

One additional but significant way to improve this process is to have the user’s preferences for information work in a more automatic and efficient way. In an earlier paper, we explained how the VIRT concept should be used to limit bits communicated to each operator to the valued and timely information⁸⁶. Specifically, information that materially affects expected outcomes should get priority, and information that is not relevant, no longer timely, nor different from what’s already believed should be automatically filtered out. This idea can be implemented by allowing operators who are concerned with various types of planned outcomes, for example, to register with an intelligent monitoring system their plans, their assumed situations and expected outcomes, and the inferences that support those expectations. The monitoring system can then take responsibility for continually reassessing the credibility of the assumed

⁸⁶ Hayes-Roth (2004), *op. cit.* See footnote 3.

situations, expected outcomes, and supporting inferences. When new information arises from any source that undercuts those beliefs, that information can be communicated to the user promptly and highlighted appropriately in that user's preferred viewer. While doing this well is an open-ended challenge, the semantic hub and associated translators constitute the essential foundation for exploiting multiple relevant sources. The hub and translators can make each source of information commensurate with each operator's specific concerns, thus enabling intelligent filtering that will increase every operator's productivity in network-centric operations.

The R&D Agenda to Achieve the Potential Benefits

So what needs to be done in research and development to implement the proposed methodology for information sharing? Because the proposed approach aims to address the useful exchange of information between all relevant suppliers and consumers, it covers a vast, open, never-closing space. For this reason, it makes sense to focus first on the highest-value problems where incremental progress can have significant impact. This explains the current focus on the *Track* concept. The inability to fuse information from diverse sources and agencies is a recognized critical weakness. Moreover, future defense and security systems aspire to creating and sharing track information on a much wider array of mobile entities than traditional military vehicles.

To advance the agenda on track-related systems, we need to accomplish several intermediate objectives, as follows:

1. Select a community of interest that recognizes the importance of this task.
2. Based on high-priority missions, enumerate and prioritize information sharing scenarios.
3. Determine a high-value near-term subset of the hub semantics.
4. Identify the viewer/editors that operators will employ in these sharing scenarios.
5. Determine the translator requirements to support the scenario.
6. Implement an initial hub and related translators.
7. Test the environment, and identify high priority requirements for improvements to the hub and translators.
8. Identify operators for whom VIRT capabilities (reduced bandwidth, intelligent filtering) have highest value.
9. Determine best methods to gain knowledge of operator's context and identify valuable information.
10. Implement query methods and notification methods to operationalize VIRT.
11. Iterate, through earlier steps, to implement continuous improvement.
12. Place responsibility for this continuous improvement process in the hands of an appropriate agent.

This proposed agenda has much in common with Department of Defense initiatives and directives. DoD has committed to using semantic meta-data to describe information in its repositories so that next-generation systems such as the Global Information Grid and Network-Centric Enterprise Services can be used to assure that each operator gets the

right information at the right time to optimize task performance. The current proposed approach enhances those initiatives by addressing the semantic requirements for enabling interoperability that improves pragmatic outcomes. In a nutshell, we've pointed out that information from multiple sources will always exist, and these will need to be inter-translated to address all operators and serve them with all sources. This approach reduces expense, risk, and "time to value" than attempting to create a single standard for representing all information. Even if it were possible, in principle, to formulate a single standard model, the pragmatic requirements would evolve faster than any standardization process could. We would forever be chasing our tails.

In contrast, this R&D agenda provides an incremental approach that can provide immediate benefits and can quickly exploit learning to gain additional benefits. A continuously improving semantic hub will be part of a *virtuous cycle*. In this cycle, users will begin to benefit from sharing some of their information. They will discover that additional benefits are potentially obtainable through broader, deeper, or more precise coverage of the semantic and pragmatic concerns. Incremental investment will yield those benefits. As the range of benefits expands, additional consumers and additional suppliers will seek to participate. The "market" for valuable, sharable information will expand. The process will feed on itself, and information sharing will be converted from an intractable problem to a continually expanding arena of exchanged value.

Going beyond Models of *Track*

It should be clear that there are many other pragmatic concerns and related concepts outside the range of those discussed here. *Track* is interesting for several reasons: (1) it's an established, if informal, concept throughout the military; (2) new threats are expanding the types of mobile entities we wish to track; (3) sharing of intelligence sources for these types of tracks is already extremely poor, so that new aspirations require a new technical foundation; (4) DOD and DHS leadership should readily understand the potential value of the *Track* as a paradigmatic focus for network-centric systems and the new emphasis on semantics and meta-information.

What other concepts might provide excellent focal points for other similar efforts? The following brief list seems easy to justify as each concept is central to some key segment of the Federal Enterprise Architecture or defense systems:

- Human Resource
- Employee
- Contractor
- Qualification
- Skill
- Knowledge Unit
- Health
- Process
- Prerequisite
- Supply
- Supplier

- Transport
- Target
- Center of Gravity
- Line of Communication
- Health maintenance
- Disability
- Repair
- Time-to-Recovery

Each of these concepts occurs in many different government agency systems, and it would prove valuable to be able to share information across such systems as well as reduce the costs of implementing similar functions in different systems. In short, wherever we have communities of interest, we'll find an overlap in pragmatic and semantic concerns. Each such overlap defines a target of opportunity for the suggested approach.

Related Research and Technology

A major emphasis in the worldwide web community and enlightened parts of government-sponsored programs these days is on the need for semantic representation and exploitation. The "semantic web" is a principal focus of the W3C. Their main objective is use well-understood ontologies (conceptual hierarchies) to annotate a wide variety of documents available on the web. These ontologies would be modeled as XML schemas and the corresponding semantic tags would be used to add "meta-data" to the text and data embedded in documents. New search methods would understand these tags and be able to make reasonable inferences about which annotated documents best address queries. DOD has already mandated that meta-data be added to describe all information bases.

Several books have appeared recently on the semantic web. One of the best is by Michael C. Daconta, Leo J. Obrst and Kevin T. Smith, entitled *The Semantic Web: A guide to the future of XML, Web Services and Knowledge Management*, Wiley, 2003. Daconta has recently taken on a leadership role in meta-data modeling for the Federal Enterprise Architecture.

Many efforts have been undertaken to develop models of relevance to military and defense applications. NIMA has the lead in DOD to develop standard geographic information models. The Open GIS Forum has created the OGIS standard for GIS abstract data types. The NATO Multilateral Interoperability Programme (MIP) has created an entity-relationship model for command-control of (mostly) ground warfare. Tracks are common presentation types in most situation assessment and command-control database-centered systems. New initiatives in DOD are aimed at creating a powerful, general Joint Track Manager. Many studies of the general fusion problem have led to a Joint Directors of Laboratories (JDL) four-level reference model describing different types of inference appropriate to combining information about military entities. This reference model has been the basis for generic architectures for fusion that could employ the type of *Track* model proposed here.

In short, there are many pre-existing attempts to use an understanding of semantics and pragmatics of tracks in various systems and applications. Our challenge is to make explicit the important elements of such work so that we can manifest it in computable semantic models. This will be a crucial step toward enabling systems of systems to interoperate and share this important kind of information in support of valued and timely decisions.

Conclusions

Many defense, homeland security, and commercial security objectives require continuous tracking of mobile entities. We wish to share information among different tracking systems working in similar domains.

To combine information from different sources, we will need a flexible framework that can tolerate and exploit data products from different systems, although these systems employ different representations and embody different assumptions. Our approach is to identify a rich semantic model of tracks that can support a wide variety of objectives related to information sharing. The semantic model is developed to play the role of a hub amidst a variety of translators. These translators implement conversions between available sources of information and the hub as well as between the hub and various viewers and editors used by human operators. In short, consumers get information that meets their needs by extracting it from relevant sources, translating it first into the hub and second into the semantic system that the consumer requires. This approach allows us to achieve significant positive benefits incrementally and offers a vastly preferable alternative to other proposed approaches.

5. Model-based Communication Networks and VIRT: Filtering Information by Value to Improve Collaborative Decision-Making

Model-based Communication Networks and VIRT: Filtering Information by Value to Improve Collaborative Decision-Making⁸⁷

Rick Hayes-Roth

hayes-roth@nps.edu

Professor, Information Sciences Department

Naval Postgraduate School

Monterey, California

April, 2005

Abstract

Command-control and other distributed, collaborative systems should achieve the best possible results with resources available. We should measure these systems in terms of the *quality* of decisions made. *Better* decisions lead to *better* outcomes, because superior choices are made about what to do, with what assets, where and when. Just as we measure manufacturing processes in terms of *value added* at each stage, we want each processing step in distributed collaborative operations to maximize the ratio of added value to cost. Both computerized agents and human personnel receive information from others, process it, and then produce additional information for others downstream in the operational processes. This paper shows that current architectures do not promote high productivity. Specifically, most current approaches encourage an increase in information supply and exchange *per se*, producing *glut* rather than *value*. This paper explains how we can significantly increase the productivity of each operator and the success of overall missions. The approach, called VIRT, treats collaborators as participants with shared models. These models determine which information is high value and for whom. The architecture gives priority to conveying high value information. Similarly, low value bits are filtered out, saving resources and optimizing value attained.

Introduction and Overview

Every modern organization aspires to improve its performance through better use of information technology. Organizations seek to increase their agility and make better, more adaptive responses to changing circumstances. As communication technology improves, organizations can operate over wider distances and can even assemble operational components on

⁸⁷ This work was supported by a research initiation grant from NPS to the author.

an ad hoc basis to meet requirements of a specific objective. People and machine-based agents operating in these organizations must collaborate with other members of their mission-oriented teams. They send information to and receive information from each of their collaborators, and each collaborative team aims to process this information to reach an effective and timely decision. After Boyd[1], the cycle of information processing and decision-making is often referred to as an OODA-loop. To be effective competitors, organizations must close these loops faster than their opponents, so that they can drive rather than react to the opponents' behaviors.

Several trends work against our ability to close our decision loops quickly. First, the number of networked collaborators is increasing, meaning that we must process information from and for an increasing number of partners. Second, the number of relevant information sources and quantity of availability data are both increasing rapidly. Third, the times available for decisions are shrinking as we seek to compete with more agile adversaries. Where cycle times in the military were traditionally anywhere from 24 to 72 hours, our aspirations are now to identify, analyze, decide, and prosecute some targets within a few minutes. Thus, making information more available and increasing its flow among collaborators ultimately reduces the quality of decisions. Just as time-sharing computers "thrash" when they become overloaded with pending tasks, people can't make good decisions when they are time-stressed and overloaded with information waiting for their attention.

The purpose of this paper is to describe a fundamental shift in the way we approach communication among time-stressed collaborators often glutted with information. The new approach is called a "model-based communication network" (or MCN). Collaborators who employ an MCN can drastically reduce the amount of work required and can significantly increase their information-processing productivity. The major benefits of MCNs are delivered through services that filter what information people and machine agents receive. We call the services that deliver *valued information at the right time* VIRT, for short. VIRT services essentially filter information so that high-valued bits are prioritized and low-valued bits are deprecated or withheld. In this way, each collaborator's incoming queue of messages is dynamically prioritized, enabling the person or agent to work on the most important information first. This "best first" approach produces the productivity gains organizations need to thrive in a networked, information-rich environment.

This paper introduces a component-based architecture for VIRT and illustrates it with examples. I describe nine VIRT components and their interconnections. After that, the paper discusses related research, current challenges and opportunities, and conclusions. This paper should prove helpful to architects, designers and implementers of systems to support network-centric operations or any environment for time-stressed collaboration and decision-making.

The Basic Problem with Current Approaches: Stateless Networking

The modern military, as other modern extended enterprises, aims (1) to exploit superb information (2) to achieve unprecedented levels of effectiveness through (3) agile, coordinated control of resources. These new enterprises form virtual organizations on an *ad hoc* basis, quickly assembling resources with needed capabilities and integrating them into a unified operational federation. To succeed, these organizations must collaboratively construct and consistently maintain a shared understanding of several important things: mission intent; alternative plans under consideration and those being executed; and the evolving situation, which includes the past, current, and future expected position and status of all relevant entities in the environment. The term common operational picture (COP) is sometimes used to mean this shared model of the battle space, but it usually connotes a more limited *view*. The term *world model*, meaning all of the facts and beliefs about the environment, is more apt. The key capability required to enable virtual organizations to coordinate and execute at maximum effectiveness in dynamic

environments is a *shared world model*. Any attempt to lay a new foundation for collaborative networks should be driven by this requirement, and putative improvements should demonstrate how they raise the quality of the shared understanding that enables synchronized, coordinated, intelligent real-time decision-making and control.

Conventional approaches to communications have focused on laying pipes that move bits using *stateless* protocols. *State* refers to what a system remembers about what it has already done and that causes it to behave differently going forward. Stateless communication is very appropriate when we are focusing on connecting mostly independent entities, for limited interactions, which arise pretty much randomly. Whatever memory is required in these interactions is supplied by the interacting entities. Usually that means two communicators begin by establishing their identities and their credentials, and then they begin to work on establishing shared state. This requires that they describe their current beliefs, identify discrepancies, and determine how to resolve those. From that point on, as long as they stay synchronized, they can quickly process new reports by a kind of triage: repetitive and redundant information can be discarded; confirming information can be coalesced into the models that “explain” it; new but unsurprising information can be accepted and used to augment to the current model; and disconfirming information can be subjected to further tasking and analysis, as appropriate. This is the maximum possible level of efficiency for handling information communicated between two parties.

Unfortunately, the actual situation in most military operations is much more complicated and much less efficient than in this idealized two-party on-going communication. Rather than maintaining continuous shared state, the communication is usually stateless. The parties don’t stay in continuous synchronized dialog. They effectively “hang up” after each short transmission. They send messages whenever they think they have information worth reporting. The longer the parties operate independently, the more their respective world models diverge. Each time a recipient receives a message, the recipient must now also attempt to determine whether differences in the senders’ world models affected what they’ve said, why they’re saying it, and how best to interpret and utilize their statements. Moreover, communications in net-centric organizations are not merely 1-to-1, but n-to-n, meaning that each party is receiving messages from others whose own states relativistically affect what they’re transmitting. Absent an absolute, common frame of reference, each communication requires the recipient to try to determine the meaning, relevance, validity, and significance for its own world model. In the fog of war, this process inevitably results in these problems: (1) many messages are sent repeatedly; (2) many recipients are overloaded; and (3) many incompatible and inconsistent views are held by different parties. The process is grossly inefficient.

The ideal communication framework for network-centric organizations is like the *blackboard* architecture[2], in which each actor can see all previous inferences and all important ideas are woven into a structure of shared beliefs. In the original blackboard architecture, the posted ideas were called *hypotheses* and these were *linked* to represent various kinds of supporting logical relationships. Publish-subscribe architectures[3] are simplified abstractions of the blackboard. In these, recipients identify what information they are interested in, and the system routes matching items from publishers to them. *Distributed blackboard* architectures are actually the best model of the ideal communication framework[4, 5]. In these, copies of subsets of the logically global blackboard are distributed to provide fast local caches for each participant, and various processes are employed to keep the replicas synchronized and consistent.

In extended and net-centric enterprises, collaborators need to share beliefs that consistently reflect their individual roles in collective plans and operations. *Plans* are an example of shared decision products best modeled as *compound objects*. These contain constituent objects that describe the elements of a plan, such as each aircraft’s mission, route, targets, refueling, weapons, etc. In addition to plans, the organizations need to share compound objects that describe their

situation analyses, including status and capabilities of blue and red forces, terrain, networks, and so forth. Each participant in intelligence, plans, and operations should be able to see a permitted view of current beliefs and should be able to make incremental adjustments to those when they have new information. Changes in beliefs should be propagated to replicas of the corresponding objects. When changes in some beliefs undercut current plans, either by nullifying some prerequisite or altering the relative desirability of a previously rejected alternative, this condition should trigger a process that reassesses the affected plans and associated analyses. By maintaining state, important news can be automatically detected in many cases, and this can allow the responsible parties to focus attention exactly where it's needed

The appropriate model for net-centric collaborative organizations should recognize their essential nature: they must be continuously synchronized though distributed, and they must be driven by significant events, those corresponding to changes that have material impacts on on-going plans. Collaborators meld and share belief structures that describe the environment, resources, capabilities, plans, and expected results of plans. These belief structures correspond to compound objects with support relationships. The communications that people should experience, because these are the ones that matter, are those that signify a material change in beliefs. Other communications should be largely invisible, as they work mostly autonomously to maintain distributed, synchronized state. Thus effective collaborative problem-solving requires: (1) the ongoing, mostly unconscious, maintenance of melded world models; and (2) the event-driven, conscious assessment of how real “news” affects previous decisions and choices for future actions. In short, models enable us to know which bits convey information because they are “news,” and we must give priority to shipping those bits to consumers who value them. Knowing how “news” changes expected outcomes for various potential consumers enables us to target the news to the right consumers promptly.

VIRT Improves Time-stressed Collaborative Decision-making

DoD has committed to transform around concepts of Information Superiority (IS) and Network Centric Operations and Warfare (NCOW)[6]. FORCEnet, as an example, aims to provide the Navy the capabilities required to support agile, rapid, precise, effective and efficient planning and operations. In these new concepts, warfighters can access and employ whatever information they need to perform their tasks. In short, every person should operate on the right information. One problem, however, is information glut. Too much information is available today, and the problem grows worse over time. Another problem is that people have to work hard to find the valuable information, either because it doesn't automatically find them or because it's buried amidst megabits of data and messages that are not important for their particular mission concerns.

Thus, IS/NCOW depends on enabling each individual to receive valuable information at the right time and, in parallel, the automatic filtering out of low-value information. This requires improved means for allowing the needs of individuals to determine just what information finds them, so they can spend more of their time assessing and acting upon high-value information. This would have the effect of increasing individual productivity throughout the military and, as a consequence, help attain the goals of IS and NCOW. Without such a capability, moreover, increasing information loads will have the paradoxical effect of reducing mission effectiveness.

To solve these problems we need a model-based communication network (MCN) that delivers to each of its customers tailored products that satisfy the objectives of “valued-information at the right-time” (VIRT). The basic VIRT method adapts the information flow around an understanding of mission plans, their rationales or *justifications*, the assumptions and forecasts they depend upon, and their expected outcomes. In short, VIRT looks for information that materially affects expected outcomes and communicates that to decision-makers so they can consider and adopt preferable alternatives in a timely way.

Here's a simple example from aviation, where a pilot's route is planned at low altitude over low mountains. The planner considers many variables, constraints, and outcome possibilities in selecting an optimum route. For one type of mission, the goal may be the shortest flight; for another it might be the stealthiest flight; and for another it might be one with best line-of-sight communications to several parties along the route. The types of variables that must be considered include: terrain elevation; winds aloft; fuel consumption and capacity; routing waypoints and their relationship to other parties in the environment; precipitation and temperature; etc. Constraints include, for example: the flight must not consume all the fuel and, in fact, the planned flight must allow for an additional hour of emergency reserve; the flight cannot encounter icing and allow for ice accumulation; the flight must maintain safe clearance above terrain, especially when winds and steep terrain interact; etc. The planner considers many alternatives for the future flight, using forecast weather data and other information and assumptions. In light of the mission objectives and assumed information, the planner chooses the best alternative for the flight plan.

Continuing with our example, considerable time usually passes between the moment when a flight is planned and when the flight actually begins. In some cases hours or even a day or more might pass. As time passes, the information available evolves and changes. Forecasts improve as their distance in the future shrinks; in addition, direct observations supply facts for what previously required assumptions. Information continues to flow into organizations like the Navy's Fleet Numerical Meteorological and Oceanographic Command (FNMOC⁸⁸) right up to the aircraft's departure and throughout the planned flight. FNMOC "knows" when weather data, for example, differ from what had been previously forecast or assumed. Enabling FNMOC to "know" which changes are significant to the pilot will then enable it to supply valued information at the right time, i.e. implement VIRT services.

In short, the supplier of information bits needs to understand its customer's sensitivities. In this example, a change from a low probability of enroute icing to a substantial probability of enroute icing would be material to the pilot. Assuming the currently preferred plan didn't violate a "no flight allowed into icing" constraint, a newsworthy violation of that constraint could arise if the forecast changes to anticipate a combination of sub-freezing temperature and visible moisture along the planned route at the planned flight time. For FNMOC to implement VIRT, it needs to know the planned route and time, constraints on acceptable flights, and a way to convey news to the operator. Each operator and plan can have its own sensitivities. Aircraft at high altitude usually are above the weather and have de-icing equipment, so they are unlikely to consider precipitation and subfreezing temperature important. On the other hand, their flight levels are more affected by "jet stream" winds and these can significantly affect fuel consumption, as just one example.

So the essence of VIRT is knowing which consumers really care about what news. Suppliers of information should monitor for a change in their information (news) that would interest operators, because it changes their beliefs about expected outcomes. The final element of VIRT consists of the conveyance employed to transmit the valued news to the user. This should include a means to highlight news in an appropriate way. Preferably, the highlighting causes recipients to attend to news with a priority that closely approximates the importance they attribute to it. Urgent and vital information deserves high priority. Unimportant data and stale information deserve low priority.

⁸⁸ I collaborated with FNMOC on the implementation of VIRT for their customers. I have been fortunate to have the knowledgeable and enthusiastic support of the former FNMOC Commander, Chris Gunderson (CAPT USN RET), and several of his talented staff, including: Bruce Gritton, FNMOC-CIO; Ensign Darin Keeter; and Doug Gentges, a contract architect/programmer. Gunderson is now Executive Director of the World Wide Consortium for the Grid (W2COG), where VIRT is a major architectural tenet. See www.w2cog.org.

We can employ a range of possible methods to implement the essence of VIRT. In the ideal world, the plans and plan evaluation methods of the operators might be known to the information suppliers. Then whenever a supplier noticed a change in relevant information, it could “simulate” the operators’ thinking to determine whether the operators would alter their previously selected plans. In just those cases, it would alert the operators. Otherwise, it would not bother to pass along insignificant changes. As a much more modest objective, we have chosen to allow operators to tell us what kinds of changes are significant to them in light of specific plans they have considered and selected. The VIRT service then takes responsibility for monitoring the identified types of changes and conveying them promptly to the interested parties.

Even this modest ambition for an initial VIRT service will produce substantial benefits. Every planner and operator is sensitive to some types of potential changes throughout the period leading up to and during plan execution. A typical pilot for a simple mission, for example, may be responsible for monitoring dozens of information types throughout a 12-hour period. An entire organization or a coordinated military mission force, for example, can make millions of “reads” to keep their plan justifications fresh. Usually, nearly all of those justification-maintaining examinations will turn up nothing valuable. As a consequence, the rare and important deviation from acceptable condition will often be overlooked.

To make VIRT and model-based communication networks routinely available, we need to provide some generic capabilities that enable suppliers and consumers of information to understand what information is valuable. These generic capabilities can then be specialized for particular domains of application and communities of practice, as when weather specialists and aviators establish a shared understanding of concerns such as “enroute icing” and “headwinds.” In the next two sections, I explain what these generic capabilities are and illustrate how we specialize them for particular applications.

Overall Technical Strategy and High-level Architecture

In this section, we consider a high-level architecture for VIRT that exploits a set of semantic models that describe the information suppliers make available to consumers. The simplified architecture is illustrated in Figure 1.

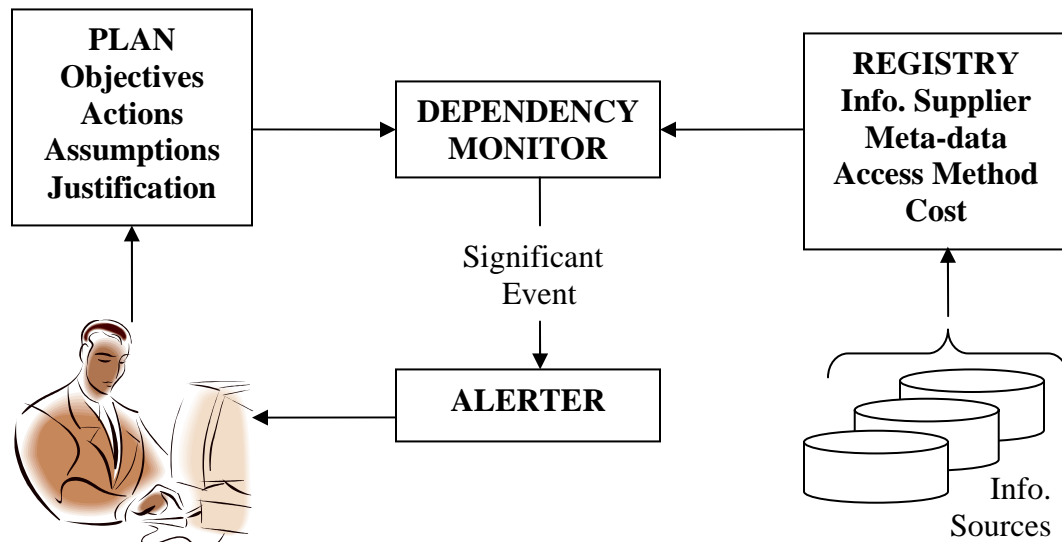


Figure 7. A simplified architecture for VIRT.

The architecture in Figure 1 simplifies much of the complexity by focusing on just a single plan, apparently planned and periodically adjusted by only the one person illustrated. Of course, real organizations comprise many teams pursuing many objectives, so there are many planners

and plans extant at any point in time. Nevertheless, the key elements of the VIRT approach appear in this simple view.

The overall flow in Figure 1 starts on the left side, where a plan has been generated. Each plan describes time-phased actions that should accomplish the plan's objectives. The planner considered what the state of the world would be at the time the plan executes, and his/her beliefs about that future state correspond to the "assumptions" recorded in the plan. When planners select a plan, they usually evaluate it and compare its costs and benefits to other alternatives. They can record their reasons for selecting one particular alternative in the form of a *justification*. A justification ordinarily explains how the planned actions should accomplish the objectives in a situation where the assumptions validly hold and, also, why the planners prefer the selected plan to the alternatives they considered. The justification often reflects that all of the alternatives considered had excessive risks or costs in comparison to the chosen one.

Let's consider a simple example. The planners might have an objective to rescue a small group of people on the ground in forested terrain. Their basic choices consist of reaching the people by ground or air and extricating them by ground or air. The likely options for ground transport include wheeled and tracked vehicles, horses, and humans. The likely options for air transport include rotorcraft v. fixed wing aircraft. Given a number of factors, they quickly reject all but the following skeletal plan:

1. Survey the area by fixed wing aircraft to find the best landing spot for a helicopter.
2. Send a helicopter with a search and rescue (SAR) team.
3. Land the helicopter at the chosen site.
4. Find and recover the party using the SAR team.
5. Depart by helicopter and return the party to a chosen medical facility.

Given this skeletal plan, the planners then focus on possible aircraft and routes, total expected flight times and associated fuel requirements, and possible time sequences for the flights. The flights become highly dependent upon the assumed wind, visibility, and icing conditions en route and at the search and rescue (SAR) area.

Let's complete the example plan. The planners assume that a 90-minute aerial survey will be required to choose the best landing site. They choose an available low-altitude aircraft that carries appropriate instruments and can reach the site with only a two-hour flight. The aircraft has adequate fuel for 6.5 hours which is just sufficient for two 2-hour legs, a 1.5 hour survey, and still leaves a 1-hour mandatory reserve. The winds in the area are forecast to be excessive between the hours of 1300 and 1800 local, and adequate sunlight is expected only from 1000 to 1900. For these reasons the flight is planned for early tomorrow morning, so that the survey begins promptly at 1000. Thus, take-off is scheduled for 0800. The helicopter is scheduled for a 3-hour flight to the search area, and is planned to depart at 0900, so that it can receive landing coordinates at 1130 from the SAR aircraft survey team 30 minutes prior to touching down.

Even this example leaves out countless details, but it provides enough to illustrate the key VIRT architecture features. The VIRT dependency monitor takes the responsibility to watch for changes in forecast or actual conditions that threaten the plan by undercutting its justification. In the current case, the monitor needs to revalidate periodically the key assumptions regarding aircraft availability, aircraft capabilities, winds, visibility, icing and fuel consumption. Table 1 shows a sample of possible changes in the world or our model of it that might undercut the plan's justification and, for each, one or more information sources that the monitor needs to reassess periodically so it can re-assure the justification.

This table lists seven out of hundreds of possible vulnerabilities of the example plan. Aircraft often have maintenance problems that ground them. If either of the planned craft are grounded before the operation is complete, the whole plan may fail. Therefore, planners must continually monitor the readiness of the craft. Similarly, the survey mission assumes the availability of various instruments, and these must continue to function until an adequate helicopter landing area is selected. Thus, mission planners should monitor and revalidate the equipment's functionality. The third item supposes that an aircraft substitution has occurred, as in response to a problem like the first or second ones discussed. In this case, the new aircraft must have as much range, load, and instrumentation capabilities as required of the one it replaced.

The fourth and fifth problems challenge the ability of the aircraft to complete the planned flights, either because the new conditions might require excessive fuel or prohibit flight. These possibilities always exist, though at varying degrees of likelihood depending on locale and season. Nevertheless, the plan is vulnerable to changes in these meteorological conditions, so the dependency monitor should continually revalidate the ability of the aircraft to fly the planned route, maintain an adequate fuel reserve, and avoid flight into icing conditions.

Table 6. Vulnerabilities and the associated dependencies monitored.

Changes that Undercut Plan Justification	Dependency that Must be Monitored
1. Planned aircraft down for maintenance	Readiness of planned aircraft
2. Survey instruments inoperative	Readiness of planned survey equipment
3. Substitute aircraft has reduced capability	Capacity and capability of replacements
4. Increased headwinds eliminate fuel reserve	Winds aloft and fuel consumption
5. Enroute icing reported by other aircraft	No icing conditions observed or forecast (temperature, precipitation, ice)
6. Survey team finds no suitable landing area	Survey objective accomplished satisfactorily
7. Departure of survey aircraft delayed	Survey objective accomplished on time

The sixth and seventh problems undercut the logic of the plan, by making it impossible for the helicopter to depart at a fixed time, receive coordinates of a landing site en route, and land shortly thereafter to perform the rescue function. The system should monitor the continuing plausibility of the assumptions the helicopter plan depends on, including the timely and satisfactory completion of the survey objective to find a landing site in time.

While there are many other ways this plan, and any plan, can be thwarted by violations of explicit or implicit assumptions[7], the point here is that machines should be employed to monitor as many important dependencies as possible. When unfolding events violate any of these key assumptions, planners should consider the consequences. The VIRT system monitors such dependencies and alerts planners when significant events occur. These significant events include the types listed in Table 2.

Table 7. Significant Events Types and Illustrative Examples.

Category of Significant Event Monitored	Illustration from Example Plan
Unavailability of required resource	Aircraft not grounded for maintenance
Inadequate capability of employed resource	Aircraft instruments operative, load adequate
Adverse change in forecast weather	Increased headwinds forecast en route
Adverse observations reported	Other pilots report high winds aloft and icing
Plan's justification negated or nullified	Resources, capabilities, and time intervals now available probably can't achieve an objective

VIRT works by seeking significant events in dynamic data sources. To do this, it must understand what significant events undercut assumptions the plan depends on and how to access and query information sources for the events of interest. The illustrative examples in Table 2, for example, might be monitored by periodically examining aircraft readiness and capability databases, wind and icing pilot reports, wind and icing updated weather forecasts for the airways and times of our intended flights, and expected start and completion times for various planned tasks that other tasks depend upon. Given a very specific list of significant events, plan dependencies, and information sources, a specialized VIRT application could be easily designed and straightforwardly implemented. We would still need to specify the best way to alert the human planners when we detected particular categories of significant events. For example, if we computed that new wind forecasts undercut the ability to maintain an adequate fuel reserve, we'd want the specialized application to consider and compute some potential workarounds, such as changes in route, altitudes, or time. The specialized VIRT application could then offer more than just "a new problem"; it could also constructively suggest a potential adaptive response. An excellent example of a specialized application prototype that addresses many of the challenges of monitoring weather for significant events and determining when and how to alert pilots is the AWARE system, described by Ruokangas and Mengshoel[8].

The simplified architecture we are describing is aimed at solving a broader generic problem, so that almost any planner can employ it for any kind of plan by seeking significant events among any pertinent information sources. Rather than a specific application, therefore, the VIRT architecture aims to provide a generic service for plan monitoring and for intelligent filtering of potentially relevant, dynamic information sources. In the generic architecture, the dependency monitor can infer what types of events to look for from any plan whose components include assumptions and a justification. The monitor can also be advised by an operator how to focus or optimize its functions. VIRT also employs a registry of available information sources to exploit whatever sources become available. VIRT is open to new information suppliers, who need only describe what their information sources are, how to access and query them, and how to reimburse or compensate the supplier if payment is required. Lastly, the architecture is open on the question of how alerts of significant events should be communicated. We expect there will be a variety of alerting methods, some more appropriate than others for particular types of users, tasks, information sources, mission objectives, significant events and equipment contexts.

In addition to the type of example plan we discussed in this section, our work on VIRT with FNMOC currently focuses on two particular Navy operational scenarios. The first of these addresses the problem of assuring that submarines receive high-value information over their limited bandwidth channels. To do that, VIRT notices when dynamically changing data differ from previously conveyed information and then determines which changes have significant import for the sub given its current mission and plans.

The second Navy scenario we are working on addresses the question of helping special operations forces minimize their risk of detection and level of effort to penetrate an enemy's defenses by minimizing their exposure to radar. In situations where the detection capability of radar depends of meteorological and oceanographic conditions, our VIRT application determines when weather and sea state change enough to justify replanning, and then triggers that replanning. In this way we hope to make SEAL missions less risky, less physically demanding, and more adaptive to dynamically changing environment parameters.

In sum, the high-level architecture of VIRT aims to improve group synchronization by understanding how changing situation variables affect their plan, monitoring specifically for potentially important changes, and rapidly alerting them to significant events that undercut the logic of their plans. We envision a VIRT system that is open to suppliers and consumers of information. The suppliers can describe what information they supply and how to access it. The consumers can describe what assumptions justify their plans and how deviations from those assumed conditions signify plan vulnerability and portend problems. Our initial VIRT projects don't try to automate the functions of re-justifying a plan in light of changing circumstances or re-planning a no-longer-appropriate plan. Those goals would require a narrow focus that could be addressed by programming a computer to solve that class of problem. Instead, we defer such ambitious problem-solving capabilities to later. This enables us to focus immediately on a generic, relatively straightforward service that can be employed by many planners, consuming many types of information sources.

Product-line, Component-based Technical Architecture

When you anticipate addressing a variety of similar application requirements with mostly generic software, the best approach is to create a component-based, product-line architecture[9, 10]. A product-line architecture defines a set of reusable generic components and specifies how data and control should flow among them to solve application problems. Over time, a successful architecture encourages developers to produce interoperable components of increasing quality and capability. Our hope is that such developments will occur to support MCN in general and VIRT services in particular. In this section, I'll describe an initial component-based architecture for these purposes.

Figure 2 illustrates the principal components of the VIRT architecture. These have some similarity to the simplified view in Figure 1, but these generic components can supply VIRT services to many different parties, related to many different plans, at the same time. Each component shown is modeled as an object with some attributes and optional methods. Generic components such as those shown can be implemented in different ways, with various specializations and enhancements. I will describe the overall component collection and interactions first at a high-level, and then do a deeper dive into each one.

We expect that VIRT services will most often be delivered in the context of organizations that plan and execute missions. A Planning Toolset component represents the types of functions and results that VIRT exploits in the planning environment. Most organizations have planning tools already, so the generic capabilities here will be obtained from existing functions augmented by some new ones. The Planning Toolset enables planners to generate candidate plans, evaluate alternatives, and justify the selections they make. Key assumptions record beliefs that a group of plans take as given. A dependency analyzer identifies particular underlying beliefs that a plan's outcome seems sensitive to. A condition generator translates these dependencies into specific conditions that should be monitored, and those conditions are monitored by corresponding Condition Monitors. Example conditions would include: "no icing en route" or "adequate fuel reserve maintained throughout plan." As time passes, elements of a plan or its assumptions may need to be updated, and a plan updater does that. This is particularly relevant as plans are executing and actual results come in to replace forecasts and expected results.

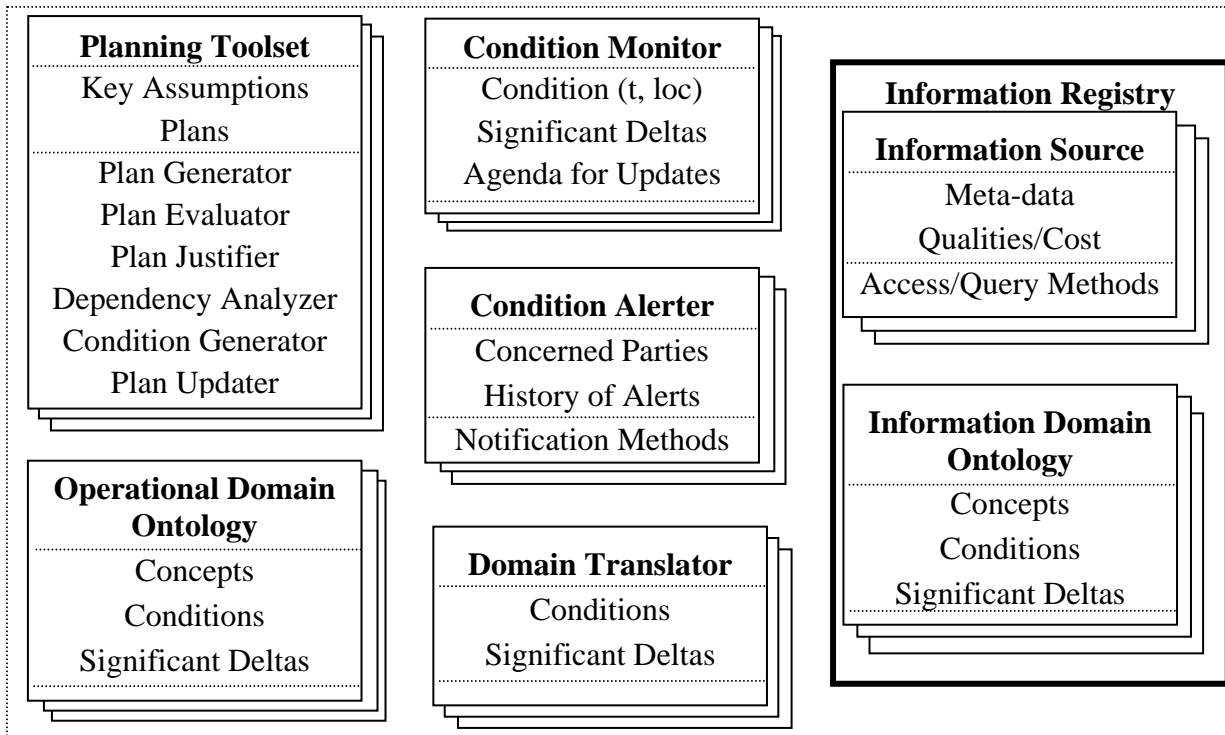


Figure 2. A component-based product-line architecture for VIRT.

Each condition generated to check and assure a plan dependency is monitored by a Condition Monitor. The condition monitor examines the value of the designated condition over appropriate time and space coordinates and records when significant changes in the value of the condition occur. It also maintains an agenda for scheduled updates to these computations. When significant events occur, as when a significant delta indicates that a condition has gone from a satisfactory to an unsatisfactory state, the associated Condition Alerter has responsibility to notify appropriate parties who are concerned with this type of alert. The alerter can use whatever notification methods the concerned parties specified for reaching them.

The Condition Monitor obtains dynamic situation data from sources described by entries in the Information Registry. Each information source provides dynamically changing data about particular variables. The variables, encodings, and other such data definitions are described in the associated meta-data. Sources may differ in terms of their perceived or rated quality and also in terms of the cost for use. Each source provides methods for accessing its data, such as particular query languages. Usually data is characterized in terms of data dictionaries or entity-relationship models.

In the future, data will be further characterized by semantic schemas or more formal ontologies that characterize what each entity and attribute means and how different values support various kinds of inference. Ontologies will be used in two very different ways. One ontology will specify the semantics of an information source, as when an attribute such as “dew point” is explained as “a 1nm x 1nm grid of temperatures at successive 3000-foot altitude bands where airborne water molecules will precipitate into visible moisture.” Another ontology will pin down the semantics the planners and operators associate with terms they’ve used in plans and conditions. For example, they may specify that “en route icing” means a condition of “sub-freezing temperature coupled with precipitation forecast in any area within 2 nm of the route occurring within 30 minutes of the planned flight through that area.”

The last component of the VIRT architecture is a Domain Translator that can translate conditions and significant deltas expressed in one ontology into a different ontology. For

example, an aviation-weather translator could translate forecast temperature and dew point information from a weather information source into “precipitation,” “sub-freezing temperature,” and “expected icing” that concern flight planners. In short, the Domain Translator relates concerns in the operational domain to data sources described in an Information Registry.

For each of these generic components, in turn, we’ll now consider their functions, interface, interconnections, and key qualities. Any particular application could then be quickly addressed by combining and specializing available implementations of these components. Each generic component is described by a table addressing the principal facets: attributes, methods, interfaces, interconnections, and qualities. An example of each component’s function is also provided. The first component considered is the Planning Toolset, in Table 3.

Table 3. Planning Toolset Generic Component Description.

FACET	VALUE	COMMENT
Name	Planning Toolset	Combines legacy planning aids with new methods for augmenting and annotating plans
Attribute	Key Assumptions	A list of conditions in the operational domain ontology underlying the plan
Attribute	Plans	A list of alternative plans, including actions to achieve the objective given key assumptions
Method	Plan Generator(obj)	Generates candidate plans to meet objective obj given key assumptions
Method	Plan Evaluator(p)	Assesses quality of plan p given key assumptions
Method	Plan Justifier(p)	Justifies why plan p achieves its objective given key assumptions
Method	Dependency Analyzer (p)	Determines how plan p results depend upon given key assumptions or other additional ones
Method	Condition Generator(p)	Generates conditions to monitor to assure plan p achieves expected results given key assumptions
Method	Plan Updater	Updates plans over time to reflect changes in key assumptions, actions or evaluation
Interface	User interface	Planners create plans, view expected results, designate conditions, specify alerts Planners receive alerts and view expectations and results
Interface	Machine interface	Creates condition monitor and condition alerter objects Provides access to key assumptions and plans
Interconnection	Uses operational ontology	Planning tools express plans in ontology terms
	Generates condition monitors & alerters	Each plan alternative has its own conditions
	Receives alerts	Each plan continually revalidated by some parties
Quality	Planning tools generate good plans	Good plans are expected to work when executed, given assumptions already made
	VIRT services easily requested with little increased workload	Requesting condition monitoring and alerting should be easy with a plan and its justification and dependencies on hand
	VIRT alerts provide	VIRT reduces bit-flow to parties by filtering out

	concise important feedback quickly	frequent redundant and immaterial data
Example Use	Flight plan created with conditions monitored for start time, icing, turbulence and fuel reserves	The plan includes a route and rate of fuel consumption; route is compared to forecast weather data translated into icing, turbulence and fuel consumption values; key conditions monitored continually as time progresses

The Planning Toolset component provides a set of functions to enable planners to formulate plans, to evaluate them, to discover and select conditions to monitor, and to update the plans as things evolve over time. The toolset component provides user interfaces to support human planners and machine interfaces to create condition monitors and alerters as well as provide access to plan and assumption attributes. The toolset component should support production of good plans and enable VIRT services to monitor important conditions without a great amount of additional work on the part of the planners.

Table 4 describes the generic component for monitoring conditions.

Table 4. Condition Monitor Generic Component Description.

FACET	VALUE	COMMENT
Name	Condition Monitor	
Attribute	Condition (t, loc)	The condition's value at time t and location loc
Attribute	Significant Deltas	Transitions (location, time, value changes) where condition value became significant
Attribute	Agenda for Updates	Schedule to get updates from info sources
Method	Update Condition(t, loc)	Using updated data, recompute condition's value
Method	Get Update	Access an info source to get updated data
Method	Accept Update	Process asynchronous data updates received from info source
Method	Identify Significant Deltas	Determine when significant changes in the condition occur
Method	Set Agenda	Determine which info sources to access and when
Interface	Machine interface	Allow planning tools to create and modify condition monitors
		Allow info sources to provide asynchronous data
Interconnection	Accesses Info Registry	Determines info sources to employ and how
	Accesses Info Sources	Accesses or queries relevant sources
	Use Domain Translator	Assesses conditions and deltas in operational domain
	Signal Condition Alserter	Notifies alserter when significant events occur
Quality	Effective at detecting significant events	Assures vigilant assessment of conditions and detection of significant deltas
	Efficient in use of costly resources	Schedules info accesses intelligently, computes as needed, and only generates significant alerts
Example Use	Change in forecast headwinds implies fuel reserve will be inadequate	The operational domain condition of "adequate fuel reserve" is computed by assuring that the difference between fuel at takeoff and fuel consumed on all route legs is enough for 60 minutes of additional flight; fuel on each route

leg is computed by multiplying average ground speed times fuel rate for that leg; ground speed is air speed minus headwind component

In the preceding table, we see how the generic Condition Monitor component works. One such object is associated with each condition. The condition is updated as new data are accessed or provided asynchronously. The monitor determines an efficient schedule for periodically accessing data sources. In the example, the condition “adequate fuel reserve” is monitored. The estimated fuel reserve is the number of minutes the aircraft can fly with fuel expected to be remaining after all planned route legs are flown. Standard parameters are used for fuel consumption per hour for various flight profiles. The key unknown is the amount of headwind. As the headwind forecast changes or actual headwind data become available for the planned route, the monitor recomputes the expected fuel consumption and then the amount remaining for reserve. If this becomes less than the required minimum, the condition goes from “green” to “red,” and a significant delta is noted. This also means a significant event, “loss of adequate fuel reserve,” must be signaled to the associated alerter.

Table 5. Condition Alerter Generic Component Description.

FACET	VALUE	COMMENT
Name	Condition Alerter	
Attribute	Concerned Parties	Identities/addresses of people or agents to notify
Attribute	History of Alerts	Record of notifications to parties
Method	Notification Methods	Means to convey alert to interested parties
Interface	Machine Interface	Receive significant events from Condition Monitor
		Access communication channels to convey alerts using notification methods
Interconnection	Accept significant events from condition monitor	Receive, note, and disseminate critical changes in conditions to concerned parties
	Convey alerts via communication channels	using notification methods
Quality	Alerts communicated promptly and concisely, avoiding annoying repetitiveness	News of significant events best conveyed to users within the context where it's most easily understood
		Repetition used only when acknowledgment required but not received
Example Use	The flight planner is notified with a short pop-up message that shifting winds have undercut the planned flight's fuel reserve requirement	After a flight is initially planned, the planner may be difficult to reach, because he or she may not be at the same computer where the plan was created; in addition to a pop-up message in an active planning window, instant messages, email, and voice messages might be appropriate; acknowledgment of any form should stop the alerting process

Table 6 describes the generic component for registering information sources.

Table 6. Information Registry Generic Component Description.

FACET	VALUE	COMMENT
Name	Information Registry	
Attribute	Information Sources	Collection of information source objects
Attribute	Information Domain Ontologies	Collection of information domain ontology objects
Method	Update Information Sources	Add, delete or modify info source objects
Method	Update Domain Ontologies	Add, delete or modify info domain ontology objects
Interface	Machine Interface	Allow access to Condition Monitor Publish updates on asynchronous channels
Interconnection	Condition Monitor reads	Monitor determines which information sources can best be used
	Domain Translator reads	Domain translator uses info domain ontology and info source meta-data to determine which data relate to the conditions being monitored
Quality	Easy to update and administer	Registry can add, delete, and change contents without limitations
	Flexibly supports diverse and evolving sources	Meta-data are described using flexible meta-meta-models, as are domain ontologies
Example Use	Registry incorporates three different sources on winds aloft, with different meta-data, and different ontologies	New sources of forecast and observed winds aloft are registered when available, and the condition monitor for “adequate fuel reserve” can employ whichever of these give best results for acceptable cost

The Information Registry described in Table 6 above provides an open architecture for new sources of potentially relevant information. Most information sources today are described, at best, in terms of the data dictionary or database schema used to store and access data. However, meta-data including semantic schemas are increasingly available utilizing XML. Moreover, standardized language systems for ontologies, such as OWL, make it likely that more semantics and domain logic will be explicitly represented and available for use. We have anticipated this trend in this architecture component. The basic role of the registry is to provide the foundation for an “open market” of information that condition monitors can access to do their work more effectively and more efficiently. Over time, suppliers should offer new and improved sources of information that planners and operate can use to monitor key conditions more accurately and cheaply.

Table 7. Information Source Generic Component Description.

COMPONENT FACET	VALUE	COMMENT
Name	Information Source	
Attribute	Meta-data	Describes data types, formats, accuracy
Attribute	Qualities/Cost	Describes source’s performance, reputation, etc. Specifies costs for use
Method	Access/Query Methods	Get dynamically updated data on request
Interface	Machine Interface	Allow Condition Monitor to read
Interconnection	Condition Monitor reads attributes	Monitor determines which data to access, when and how

Quality	Condition Monitor accesses & queries	Monitor uses source's methods to obtain the desired data
	Asynchronous data updates to Monitor	Source uses appropriate channel to convey data updates asynchronously
	Meta-data accurate	Contents consistent with descriptions
	Qualities/costs accurate	Performance consistent with descriptions
Example Use	Access/query methods reliable and produce concise results	Methods work as advertised and don't produce extensive amounts of extraneous bits that must be processed
	Winds aloft source	Winds aloft relevant to a route are produced by NOAA; query gives a table of wind direction and speed, along with temperature, at the time of flight, at each altitude that is a multiple of 3000', updated twice per day, reported by major air traffic regions

Table 7 describes the generic component for an Information Source. Each information source provides an independent set of data appropriate to various concerns. Typically sources correspond to periodic products of organizations such as NOAA and FNMOC for weather, and military, financial, commercial, maritime and various other products from corresponding organizations. Each Information Source describes its own data using meta-data techniques like those of XML or OMG's Model Driven Architecture (MDA). The source advertises its quality and costs, such as its reputation with consumers for timely, accurate, reliable performance. It provides ways to query and access its data. This may include techniques for posting "standing queries" that cause the source to transmit new, relevant information asynchronously to the requestor. In the example, a typical source for winds aloft is used by the monitor for the "adequate fuel reserve" condition. This source provides the wind direction and speed in a broad area at various elevations. The fuel reserve condition would use estimated headwinds by flight phase and route segment. It might employ piecewise linear approximations based on wind forecasts from different air traffic centers the route crosses. It could also interpolate between forecast altitudes as required to match a planned flight altitude. The example might have shown another information source that could provide more precise headwind estimates or, perhaps, an accurate fuel consumption estimate based on detailed wind modeling if those were available. When new sources become available, the architecture aims to make it easy to exploit them rapidly.

Table 8. Information Domain Ontology Generic Component Description.

FACET	VALUE	COMMENT
Name	Information Domain Ontology	
Attribute	Concepts	Terms used and their semantic properties
Attribute	Conditions	Propositions and operators used to specify important situational characteristics
Attribute	Significant Deltas	Minimum changes in calculated conditions worthy of attention
Interface	Machine Interface	Condition Monitor may read attributes Domain Translator may read attributes
Interconnection	Condition Monitor reads attributes Domain Translator reads attributes	Monitor determines which concepts, conditions and deltas pertain to its tasks Translator maps concepts and conditions from information domain to the operational domain
Quality	Ontology expressive and interpretable	Language used should simplify writing/editing Machines can easily perform required inferences

	Deltas as small as necessary and as large as permissible	Appropriate deltas reduce workload
Example Use	Winds aloft concepts, part of aviation ontology	Winds aloft described as a dynamic vector field over 3D space above surface; at each point, the variable has an amplitude in knots and a direction given with respect to true North; field is approximated by a grid, comprising spatial regions associated with air traffic centers and altitude levels, in multiples of 3000' above mean sea level; forecast values are valid for six hour intervals; forecasts updated twice per day

Table 8 describes the generic component for an Information Domain Ontology. Ontologies have been used in computing for more than a decade, but only recently have they become commonplace. An ontology is a description of the semantic concepts of a domain, including important relations among those concepts. The simplest, most familiar ontologies come from biology, where Linnaean taxonomies describe plant and animal class relationships. Ontologies reduce the complexity of organizing facts. They also economize the recording of inferable facts. For example, we know all mammals have fur and nursing females lactate; therefore we can infer that our own female dog will lactate when she gives birth to puppies.

Our architectural component indicates that three attributes will be most important. We want to know the concepts addressed by an information source as well as the conditions it addresses. In the example, we can see how wind velocity and direction can be culled out of the data grid. With additional ontology definitions, we can relate these conditions to others of interest, such as headwind component along a route segment and, ultimately, the fuel consumed and the reserve remaining. Standards for ontology representations are emerging, and we expect the information domain ontologies will become increasingly standardized. Prior to becoming standardized, however, we should expect that ontologies will evolve through use and experience among a community of practice. Each important operational problem requires information suppliers to meet the needs of planners and operators. This means that the ontologies will become increasingly adapted for use in condition monitoring and translation. The value of information, in short, derives from its ability to materially improve expected outcomes of operators' plans. Important and useful distinctions find their way into the concepts and conditions of the ontology, and these in turn determine the data that the suppliers report in their information sources.

Table 9. Operational Domain Ontology Generic Component Description.

FACET	VALUE	COMMENT
Name	Operational Domain Ontology	Same structure as Information Domain Ontology, but reflects operator concerns
Attribute	Concepts	Terms used and their semantic properties
Attribute	Conditions	Propositions and operators used to specify important situational characteristics
Attribute	Significant Deltas	Minimum changes in calculated conditions worthy of attention
Interface	Machine Interface	Planning Toolset can read attributes Condition Monitor can read attributes Domain Translator can read attributes
Interconnection	Planning Toolset reads attributes Condition Monitor reads	Planning tools use operational concepts to specify plans, assumptions, conditions Monitor determines which concepts, conditions

	attributes	and deltas pertain to its tasks
	Domain Translator reads attributes	Translator maps concepts and conditions from information domain to the operational domain
Quality	Ontology expressive and interpretable	Language used should simplify writing/editing
	Deltas as small as necessary and as large as permissible	Machines can easily perform required inferences
		Appropriate deltas reduce workload
Example Use	Winds aloft concepts, part of aviation ontology	Winds aloft described in terms of headwind and tailwind conditions along planned route of flight at planned time of flight; these decrement or increment airspeed to produce estimated groundspeed; groundspeed determines elapsed time for each route segment
	Take-off fuel quantity	
	Fuel consumption	
	Adequate fuel reserve	

Table 9 describes the generic component for the Operational Domain Ontology. This ontology is entirely similar to the ontology for the information domain, but it describes directly the concerns planners and operators have, rather than using the terms and codes of information suppliers. In the example, winds aloft are conceived in terms of their impact on groundspeed, flight time, fuel consumption, and the concern for adequate fuel reserve. Most operators today do the translation from information sources into operational domain concepts in their heads, routinely, often many times per day. The VIRT architecture addresses the need to off-load such computation onto machines and to make it be “exception-driven” rather than intensive, repetitive, and usually immaterial. As the operational communities learn the value of making their concerns explicit, the planning tools will evolve to use the ontology concepts and conditions for human interface with the operators. In addition, dependencies will be converted into key conditions for monitoring. Lastly, the operational domain ontology will define the target range of the domain translator that can map source information into operational concerns.

Table 10. Domain Translator Generic Component Description.

FACET	VALUE	COMMENT
Name	Domain Translator	
Attribute	Conditions	Conditions the translator can infer in the target ontology
Attribute	Significant Deltas	Deltas the translator can infer in the target ontology
Interface	Machine Interface	Condition Monitor may employ translator
Interconnection	Condition Monitor employs	Condition Monitor evaluates operational conditions in part by inferring their values as translations of computed information domain values
Quality	Coverage	Translations available for important conditions
	Efficiency	Machines can easily perform required inferences
	Correctness	Translations and inferred values not erroneous
Example Use	Headwinds and tailwinds inferred from planned route, time, and winds aloft	The tailwind for each route segment is computed by finding the appropriate forecast wind aloft vector and computing the component parallel to the direction of flight; the headwind is the negative of the inferred tailwind

Table 10 describes the generic component for translating beliefs and values in one domain ontology into another. In many important applications, this is straightforward. Computations can

be arranged either as goal-driven or change-driven. Goal-driven programs are asked to determine some values of a parameterized description, such as Headwind(route-segment-1, ?speed?). The interpreter of the ontology mapping then determines the actual value for the parameter ?speed? and returns an assertion with the parameter replaced by the actual speed along route-segment-1. Data-driven programs respond to changed observations and propagate inferences to the parties who have indicated a continuing need to be informed. Thus, when the twice-daily winds-aloft forecast is published, the headwind along each segment of each plan could be recomputed to determine if any significant deltas occurred. In that case, the new headwind value for the affected route segment would be conveyed to the appropriate parties.

Translation between formal languages has been a focus of computing research for decades. There are many simple to use language systems that can be used to build specialized translators. General-purpose, generic translators can also be built for a wide range of descriptive ontologies. The most general form of the translation problem is, in principle, not solvable, however. But that limitation isn't expected to have any practical impact on most applications of VIRT services, because these are likely to address practical domains where operators already do translation of this sort routinely, usually in their heads. Automating that work and doing it systematically for important conditions should produce significant value for planners and operators.

This completes the current description of the VIRT product-line architecture. We do not yet have much experience with actual implementations, and no off-the-shelf implementations exist for the components. We are, thus, at the start of what should be a long-term, fruitful cycle of evolution, continuous improvement, and architectural refinement. The goal of framing an architecture at the outset is to encourage an approach that favors openness, reuse of assets, and a focus on quality components. These should make the benefits of VIRT available to more people, sooner, at lower cost.

Related Research

Many people have touched on aspects of model-based communication networks, adaptive replanning, information filtering, and selective information push. The principal related research is summarized briefly here.

Psychologists, sociologists, and students of decision-making and communication have demonstrated the importance of shared beliefs and shared context in interpersonal dialog and collaboration. We are all familiar with the phenomenon of communications becoming briefer among teammates, family members, and colleagues as familiarity and experience increases. In Shannon's original treatise on information theory, he characterized a single *bit* as the amount of information required to reduce 50% of the receiver's uncertainty. This means that the more communicating partners share beliefs, the less uncertainty they have, and the fewer bits they need to specify a preferred option.

In military and business operations, planners work to achieve similar communication efficiencies. They do this in multiple ways. First, they adopt specialized terms to characterize problem contexts, relevant potential actions and resources, and criteria by which possible plans should be judged. In addition to defining concepts embodying the key distinctions that clarify choices, they often adopt short-hand jargon. The specialized language and methods of communities of practice has recently been recognized as an important foundation for building effective systems.

Much of the shared context that simplifies communication between collaborators often consists of the perceived situation or its externalized representation. In the military, for example, we wish to provide all collaborators access to a *common operational picture (COP)* that portrays

the battle space, actors, behaviors and intentions. In some human activities, the externalized representations have high “ecological validity.” For example, in manufacturing, CAD/CAM technologies nearly guarantee that the “drawing” *is* the “part.” In mountaineering and war, however, the “map is *not* the terrain.” A COP is a constructed, compound, complex hypothesis. It never corresponds exactly to reality. Nevertheless, when teammates can see and share the same externalized representation, they can significantly reduce the volume of bits they exchange in order to designate an entity or characterize an option.

Several trends are pushing people beyond the point at which they can perform adequately in these contexts. First, the volume of potentially relevant information is increasing exponentially. Second, the required cycle times for adaptive response are shrinking. And third, the teammates increasingly are geographically distributed, often culturally dissimilar, and unfamiliar with one another. In these contexts, the informal methods that have enabled people to exchange a few bits in a face-to-face interaction with familiar colleagues won’t suffice. Large portions of their collaborative work will have to be somewhat formalized so that computers can perform an increasing proportion of the information processing required.

Research on “human-machine symbiosis” traces its heritage to a famous paper by Licklider, who was the director of ARPA’s Information Processing Techniques Office (IPTO). His vision continues alive today, and has already taken us through time sharing, personal computing with graphical user interfaces, the Internet and now the Web. Yet none of these developments have succeeded in enabling machines to off-load a great deal of the information filtering appropriate for planners and operators. This requires information consumers to clarify what they need to know and information suppliers to clarify what they provide. It then requires the machine to help translate from the source ontologies of suppliers into the operational ontologies of consumers. It’s my belief that the foundations exist for all the capabilities presupposed in the architecture, but a concentrated effort needs to be undertaken to implement these and refine them to the point that communities of practice can regularly employ them. Allowing operators and suppliers to “close the loop” is critical: ontologies, translations, and key conditions will all need to be refined through experience. Thus, all the tools need to be in the hands of the producers and consumers of information. Just as spreadsheet programs launched a revolution in business modeling and business use of interactive computing, I anticipate MCN and VIRT will launch a revolution in military and civilian use of ontologies and model-based information filtering for collaborative decision-making.

Principal Remaining Challenges

There are three principal challenges to making MCNs ubiquitous. First, we need practical ontologies for important domains. Second, we need leading operational communities to transform their processes around the “management-by-exception” style of VIRT. Third, we need open, evolutionary markets for information suppliers and consumers. While each of these has technological aspects, the more challenging aspects of each concern the managerial approach taken to business processes and the required transformations.

Before PCs and the rise of the Web, computer users expected to obtain systems from professional programmers. Systems were specified and acquired through sizable and difficult contracting arrangements. Some successful systems were procured this way. Many procurements failed. The major causes of failure were delay, cost overrun, and lack of usability or effectiveness. Simply stated, what computer users want is difficult to state, is often unknown to them, and usually changes over the course of months or years. Procurements are thus shooting at an ill-defined and moving target in many cases. Missing the target, then, should not be surprising.

The PC, with its personal applications such as spreadsheets, and the Web, with its ubiquitous authoring, hosting, and editing tools, have created a vibrant community of information suppliers

and consumers. Many suppliers are professionals associated with businesses. Many suppliers are individuals. The trend is moving toward more suppliers, creating more sources, updating them more often: in short, more sources, more dynamism.

Organizations need to transform around the potential of dynamic information, agile response, distributed virtual enterprises, and self-synchronization, as in NCOW/IS. This means organizations must engage in continuous evolution, shifting their processes increasingly around better uses of information and computing. The “learning organization” is an excellent description of the new “business as usual” enterprise. MCNs enable the far-flung partners in an enterprise to collaborate succinctly, relying on externalized representations of common beliefs to eliminate the need for much redundant information exchange. VIRT services enable each planner and operator to off-load responsibility for continuous, repetitive review of important conditions to machines. All of this depends only on the development and continuous improvement of the ontologies, the shared models of how important things work in the domains of interest. Finally, an information market will enable new suppliers to proffer innovative information sources that VIRT condition monitors can automatically access. As part of such a market, whether commercial or controlled, sources need to be rated for quality and cost, so that consumers can exploit the most advantageous ones.

Organizations such as Navy FNMOC understand the importance of transforming from a traditional supplier of commodity information products to a key partner of operators who value important and timely information. FNMOC provides an excellent example of the leadership required to move into an important role in the network-centric future.

Near-term Exploitation Opportunities

I have identified several good opportunities for near-term exploitation of the VIRT services, including several with FNMOC as early collaborative partners in this effort. In essence, weather and oceanographic data are important to most soldiers, sailors and aviators. They examine copious amounts of data when generating plans, continually revalidating plans, and conducting operations. Many of these operations can achieve better outcomes with reduced risk if they can receive and exploit improved forecasts or more accurate, timely updates. However, operators can’t spend all their time looking at streams of dynamically updated weather data. For them to exploit the advantage of more and better information, they need vastly improved and automated filtering. VIRT is being applied to a number of applications within FNMOC. As one example, we aim to reduce the probability of detection of stealthy missions by assuring that the planners and operators receive valuable information at the right time.

As Ruokangas and Mengshoel demonstrated with their AWARE prototype, almost everyone in the aviation business can benefit from automated filtering and condition monitoring. Everyone who plans routes that are subject to unpleasant surprises would benefit from VIRT monitoring.

The DoD hopes to provide every level of command-control decision-making improved tools for situation awareness and real-time agile response. The COP is a foundation of this vision. The COP should be reconceptualized as a composite hypothesis of constituent models of the battle space and actors. Each component model, such as a blue-force flight or a neutral ship, can autonomously update its own expected state consistent with our knowledge of its plans and normal behavior. This obviates communication of unsurprising state changes. This allows the communication volume to be reduced to just those bits of *information*, corresponding to surprises or reductions in uncertainty. In this way, distributed collaborators can achieve a higher level of shared understanding with reduced volumes of communication. This, in turn, means they can spend more time on high-value activities, rather than being kept busy processing low-value data.

Organizations that exist to supply important information to planners and operators have a great opportunity to begin moving into the new paradigm that the VIRT architecture describes. Their products will be more valued if they are characterized by ontologies and if these are related to and translatable into operational domain ontologies. In the case of weather, as an example, the most successful organizations should be measuring their progress in terms of the import, ease, efficiency, and timeliness planners and operators attribute to their products. These are the kinds of ratings that will become advertisements and evaluation criteria as the open information market develops. Being first and best at the new game can establish significant, potentially permanent, competitive advantages and leadership.

In short, the best opportunities will arise with the organizations most eager to accelerate the transformation into net-centric, information-superior enterprises. As with other generic technologies, the real question isn't whether VIRT is applicable, but "Who's ready now?"

Summary and Conclusion

This paper has described a vision of a transformed way of operating, especially for organizations that routinely plan and execute plans. The need for this transformation arises from both new problems and new opportunities. The new problems concern new kinds of competitive challenges and new pressures to behave with greater speed, agility, and precision. As the types and volume of potentially relevant information increase without bounds, the pressures on humans to produce excellent decisions and outcomes become unrealistic. Humans need to exploit computing power to reduce their tasks to a manageable level. For our organizations to get the best results, the human resources need to spend their limited time on the most important things. MCN and VIRT provide frameworks for doing that. This new architecture exploits several significant opportunities that have been developed over the last decade: (1) networked communication; (2) ontologies and inference; (3) information filtering by machines; and (4) incredibly cheap computers.

The architecture proposed here must be implemented in specialized applications with particular supplier and operator communities to prove its worth and thus become "obvious" to a larger population. As with many new "obvious" technologies, the early successes require leadership and pioneering experiments. Some of this is now underway, but much more needs to be done. The point of this paper is to provide a simple trail map for pioneers to follow.

References

1. Coram, R., *Boyd: The fighter pilot who changed the art of war*. 2002, Boston: Little, Brown.
2. Erman, L.D., et al., *The Hearsay-II Speech-understanding system: Integrating knowledge to resolve uncertainty*. Computing Surveys, 1980. **12**(2).
3. Group, O.R.-T.S., *Data Distribution Service for Real-Time Systems Specification*, in *OMG Adopted Specification*. 2003, Object Mgt. Group.
4. Lesser, V.R. and D.D. Corkill, *Functionally-accurate, cooperative distributed systems*. IEEE Transactions on Systems, Man, and Cybernetics, 1981. **SMC-11**: p. 81-96.
5. Wesson, R., et al., *Network structures for distributed situation assessment*, in *Readings in Distributed Artificial Intelligence*, A. Bond and L. Gasser, Editors. 1981, Morgan Kaufmann. p. 71-89.
6. Alberts, D.S., J.J. Garstka, and F.P. Stein, *Network Centric Warfare: Developing and Leveraging Information Superiority*. 2nd ed. 2002, Washington, D.C.: Office of the Secretary of Defense (ASD/C3I/CCRP).

7. Hayes-Roth, F., *Using proofs and refutations to learn from experience*, in *Machine Learning*, R.S. Michalski, J.G. Carbonell, and T.M. Mitchell, Editors. 1983, Tioga Publishing: Palo Alto, CA. p. 221-240.
8. Ruokangas, C.C. and O.L. Mengshoel. *Information filtering using bayesian networks: effective user interfaces for aviation weather data*. in *Proceedings of the 8th international conference on Intelligent user interfaces*. 2003: ACM.
9. Bosch, J., *Design and use of software architectures*. 2000: Addison-Wesley.
10. Clements, P., R. Kazman, and M. Klein, *Evaluating software architectures: methods and case studies*. 2002: Addison-Wesley.

6. Two Theories of Process Design for Information Superiority: *Smart Pull* vs. *Smart Push*

Title:

**Two Theories of Process Design for Information Superiority:
Smart Pull vs. Smart Push**

Topics (in descending order of apparent relevance):

C2 Architecture, Policy, Network-Centric Metrics, C2 Concepts and Organizations, C2 Analysis, Cognitive Domain Issues

Author information:

Rick Hayes-Roth

Professor, Information Sciences Dept., Naval Postgraduate School

email: hayes-roth@nps.edu

589 Dyer Road, Bldg. 235

Root Hall 223, Monterey, CA 93943

Telephone: 831-656-3983 or 650-327-1166

Fax: 208-445-0127

Abstract

This paper asks how information should flow among networked entities in NCOW. In particular, should the entities *actively seek*, acquire and process relevant information or should they *wait to react* to information that others send to them? In short, should they *pull* information or should they rely upon others to *push* information to them? In most tactical contexts, *smart push* will improve efficiency by orders of magnitude compared to *smart pull*. Our analysis reveals that efficient information processing chains require a general capability to watch for key events. Humans and the computer applications supporting them will use this capability to detect events matching *conditions of interest* they specify. This capability plays a key role in transforming networks into integrated value chains. Where traditional networks aim at supporting unregulated exchanges for data bit flows best suited to random access and unpredictable process sequences, the capability to delegate condition monitoring enables us to transform networks into conveyers of *timely, valuable information*. To maximize efficiency, we must use processes where each successive step receives just such valuable information as its input. Thus, condition monitoring and its associated smart push constitute a required foundation for the efficient process chains needed to achieve information superiority.

Background

Two basic alternatives exist for providing needed inputs to process steps, whether we are discussing supply chains of material goods or information processing chains associated with decision-making and control. Processing entities⁸⁹ can seek out relevant inputs and, upon finding them, procure them. Or the entities can inform suppliers of their requirements and depend on the suppliers to deliver needed inputs at the right time and place. A major shift occurred in manufacturing over the last two decades associated with “just in time” (JIT) approaches and “supply chain integration.” Top-performing enterprises shifted from the former process design approach to the latter one. When suppliers have an excellent understanding of their customers’ processes, schedules, and ongoing process state, they can deliver valued inputs just when they are needed. In manufacturing, this reduces costs in significant ways, especially by reducing inventory and work in process (WIP) that consume space, time, and processing resources associated with storing, searching, and moving around those items not immediately required by the next processing stage. Thus, inappropriate or low quality inputs, as well as inputs received at an inopportune time, increase costs and may actually represent *negative* value.

Systems designed to produce quality decisions have many parallels with manufacturing systems, though the former work by adding value to “bits” where the latter add value by transforming “molecules.” The key to high performance, in both cases, is to produce the most valued products as efficiently as possible. Value reflects the degree to which the products embody superior features and qualities and get to “market” promptly. Efficiency, on the other hand, means producing these valued products with a minimal consumption of resources. To achieve that efficiency we use the best tools within well-designed processes. Usually, we improve efficiency by eliminating as many sources of friction as possible that consume resources unnecessarily or increase latency. Friction may be physical or virtual, as when differences in information systems introduce delays and difficulties in accomplishing successive steps.

In moving to network-centric operations and warfare (NCOW), the Department of Defense (DoD) has recognized the importance of reducing friction that makes sharing of information between different entities difficult. When information isn’t represented in a standardized way, an entity that wants to find it, procure it, and apply it can incur major delays and difficulties. So, DoD is moving to make all information readily locatable, readable, and interpretable (Wolfowitz, 2004). It hopes to establish consensual meta-data schemas to accomplish this. Even if this approach works, it leaves the question of information *logistics* still unanswered: how can the information supply chain be integrated most efficiently? In particular, should decision-makers seek out and *pull* the information they need, or should suppliers *push* the right information to them at the right place and at the right time?

The current approach to the design of the DoD Global Information Grid (GIG) and its Network-Centric Enterprise Services (NCES) emphasizes *pull*. The work being done by my colleagues at the World-Wide Consortium for the Grid (W2COG) and me, on the other hand, focuses on the alternative approach. In our approach, networks are designed

⁸⁹ Some entities that make decisions and perform actions are human and others are information-processing machines. When the military performers are humans, we often call them *operators*. Frequently, we’ll just use the abstract term “processing entity” to ignore the nature of the process performer.

to optimize information logistics by implementing what we call *Valuable Information at the Right Time* (VIRT). In this approach, suppliers work with intelligent computing machinery to determine which bits should flow to which consumers, thereby integrating the information supply chain in a manner parallel to the recent advances in manufacturing.

This paper aims to clarify the two alternatives and expose the conditions under which each provides a superior information logistics solution.

Proposed Approaches

The information logistics⁹⁰ problem is concerned with optimizing the flow of bits to processing entities distributed across a network. We could formalize this, but it's probably most useful to provide an informal, intuitive characterization of the problem. We consider any number of information producers, who can operate on various inputs to produce information outputs and, in addition, any number of decision-makers that convert selected inputs into outputs that correspond to choices. Some of those choices might trigger actions, through coupling to effectors. Other choices become information inputs to additional processing entities. Some of the entities are people and some are machines or software programs running on machines.

In all systems various limitations constrain attainable results. In most of the distributed systems we are concerned with, constraints limit how much can be done, how well, and how quickly. If we are trying to conduct a battle or minimize casualties from a natural disaster, we typically have too few people, too little time, too little information, and too little available computing resources. Different systems in different contexts will experience these constraints in different orders, but the constraint induced by limited human processing resources typifies crisis response situations. The information logistics challenge is to optimize the quality of results obtained by such a system when resource constraints limit its effectiveness. The essential question facing system designers is how should we select and sequence information for each processing entity to maximize the value of our results?

Very complex systems can't be analyzed and optimized using a closed form approach. We must rely upon heuristics to guide our design in such cases. One obvious heuristic to embrace in systems of the sort being considered is as follows: "Whenever possible, favor behaviors that make better decisions faster." Equivalently, we want to design processes that improve the quality of decisions and increase the speed of decision-making.

For our analysis, we suppose that valued results are produced through the application of systematic processes. For example, manufacturers make many different mixes and different products as a result of each product instance flowing through a set of process steps specified for that type of product. In a similar way, we view information processing systems as producing information products. These products include decisions and associated actions, resulting from a series of process steps applied to appropriate inputs. A number of steps are performed by people, and others are accomplished by computerized agents. Regardless of the type of processing entity, each step transforms

⁹⁰ *Distributed information logistics* is introduced and explained in Chow, *et al.* (2000).

input information into output information. The quality of the output can be assessed along various attributes. The single best measure, where we can obtain it, would be the *value* of the information. The value of an interim result corresponds to the expected improvement in eventual final products attributable to it. If final results improve because some interim decision was reached, the increase attributable to that decision is its value. While difficult to measure precisely, estimation of *value added* is a routine part of all mature supply chains. Processing steps and interim results that don't improve expected outcomes may have zero or negative value, because they consume resources without producing benefit.

In short, information logistics addresses the question of how information should flow among processing entities to optimize the value attained. While no definitive optimum may be obtainable, heuristics employed in supply chain integration seem applicable to the flow of information in NCOW chains. Some of these heuristics are listed below:

Desirable Behaviors

- Keep your most specialized, expensive processing entities busy
- Minimize lateness on the most valuable, time-sensitive products
- Drop low-value products before sacrificing high-value ones
- Minimize product quality problems that ensue from low-quality inputs
- Minimize time lost to set-up and cut-over required to handle changes in products or changes in inputs

Undesirable Behaviors

- Load your processing entities with more inputs than they can process
- Allow some of your required processing entities to wait for needed inputs
- Make processing entities filter out low-quality inputs
- Make processing entities prioritize backlogged tasks
- Make processing entities adapt to input variability

Thus, the basic strategy in achieving optimum product output is to allocate as much of your resources as possible to adding value, working on the highest value tasks where possible. Actions or interim results that waste time, waste processing resources, or produce little or no value should be avoided. While this may not achieve a provable optimum result, systems that adhere to these heuristics clearly outperform those that do not. As a consequence, we will prefer the *better* systems to the *poorer* ones, moving up so-called Pareto frontiers as high as possible.

In the simplest, most extreme form, there are two alternative approaches to the question of managing information flow in these distributed NCOW systems:

- **Theory 1:** Describe all information available using some type of meta-data description. Give each processing entity good search tools. Have each entity seek and acquire whatever information it needs, when and as needed.
- **Theory 2:** Have each processing entity describe conditions that would make its current plans undesirable, because those conditions would contradict assumptions needed to justify the choice of the affected plan. Enable agents to alert the affected entity. Enable the alerted entity to respond quickly to the received news.

In my classes, I have referred to these two approaches simply as **Theory 1** and **Theory 2**. Here, however, we might benefit from recognizing that the first is actually a “pull” approach to information logistics, whereas the second is a “push” approach. In each case, when the tools and computing software exploit domain semantics, ontologies, rules of inference and similar elements of artificial intelligence to improve understanding and performance, we can readily refer to these as “smart,” “intelligent” or “knowledge-based.” So when an operator in a Theory 1 system can ask for all recent reports about the geographic quadrant he currently occupies, the query processor can employ much knowledge and intelligent reasoning. Similarly, if an operator posts a planned route of flight, altitude, and expected waypoint times, an agent that filters spatially and temporally relevant pilot reports would illustrate a degree of intelligent push.

So both approaches can simplify the work imposed upon a human operator by enabling that operator to express queries in terms of high-level semantic concepts and domain-relevant conditions of interest. As a consequence, our analysis will mostly ignore any questions of whether or how knowledge-based intelligent processing might be performed. Another path we won’t explore concerns various ways to mix the two approaches. Instead, we want to focus on the relevant strengths and weaknesses of pull v. push approaches to the information logistics question. Our purpose is to develop a crisp appreciation of the two pure and opposite management strategies.

Analysis

Theory 1 is mostly consistent with current US architecture and management directives. Its proponents have been motivated by several observations. First, the US has notorious difficulty sharing information in a timely way, especially from intelligence sources to diverse military and homeland defense operators (National Commission, 2004). Significant backlogs of intelligence observations often sit for long periods before analysts can work their way through them. As a result, intelligence data often don’t get to people while they’re still current and valuable. The former DoD CIO, John Stenbit, recommended a new information logistics process that posts data as soon as they’re available, even before analysts could interpret them and convert “data” into “information” (Ackerman, 2003). Further, former Deputy Secretary of Defense Paul Wolfowitz issued a directive mandating that all DOD entities possessing potentially useful data should describe those data using an appropriate XML meta-data schema as a first step toward enabling all operators to find and access those data (Wolfowitz, 2004). These steps have been leveraged by the DoD National Information Infrastructure (NII) plans for the Global Information Grid and Network-Centric Enterprise Services (Stenbit, 2004). In their thinking, information superiority will be achieved by providing excellent search tools so that each operator can find and acquire information needed.

Theory 2 was implicit in some of my earlier papers, such as [Hayes-Roth, 2004a, 2004b, 2005a]. It represents a strategy for exploiting the exponential increase in the processing capacity of machines to enable humans with fixed, finite capacities to benefit from increasingly vast quantities of relevant data. Theory 2 aims to mitigate what humans experience as the “data glut.” Seeing the similarity between manufacturing supply chains and information decision chains also motivates us to ask how we can emulate the

efficiencies of just-in-time processes with lean, low-inventory processes. The challenge in achieving this kind of low-friction, efficient integration in manufacturing has been to make suppliers adapt their products and methods to optimize their customers' performance. This requires that suppliers get smart about how their customers use the suppliers' products and how users' costs can be minimized. As a simple example, FedEx discovered that customers needed to schedule and track shipping from their own premises, so FedEx moved those capabilities to the customers' workstations and adapted their own systems, schedules, and services so the customers themselves could view and control them. Thus, while shipping services were outsourced, each customer actually brought more control of adapted shipping capabilities into their own processes.

The VIRT methodology assumes information supply chains are restructured in a manner akin to such manufacturing supply chain adaptations. In particular, it assumes that suppliers of information learn what customers actually need to know and provide just that information to them. Intelligent agents, playing the role of information brokers, accept statements from processing entities or operators describing "conditions of interest" (COIs). These conditions describe potential events that would motivate the operators to change their planned actions to achieve better outcomes.

As an example, consider a helicopter pilot intending to fly a particular route in hostile territory. During planning, the pilot plots a low risk route. Subsequently, new observations about anti-aircraft emplacements along the planned route of flight would match one of the pilot's COIs.

In short, Theory 2 assumes many processing entities have a continuous need to know about things that undercut their previous decisions by violating some assumptions on which those decisions depend. Operators engaged in real-time plan execution consider such information *vital* and *urgent*. Those attributes, significance and timeliness, make the information high-value. Theory 2 sees information logistics as the challenge of (1) getting high-value information quickly to operators who are dependent on it and (2) assuring operators give priority to received high-value information so they adapt and achieve better outcomes.

Figures 1 and 2 below diagram simple functional models of the value chains associated with the two different types of processing organizations. These models use a lot of informal shorthand to make them simple and easy to read. The first model is centered on the Processing Entities PE_1, \dots, PE_k that do work. The PEs add value by accessing various information sources IS_1, \dots, IS_n to produce valued products labeled v . Each PE acquires its inputs through interaction with a Query Specifier (QS). The function q on the link between the PE and QS represents the transaction that yields information products from QS with various levels of efficiency. So, for example, we can assume q gives a lot of information at low cost, as most query processing systems do. Google, for example, gives thousands of relevant answers to most queries.

The rest of the process works roughly as follows. Once a query is specified, the transaction p translates the query into a query plan by working with the Query Planner (QP). The transaction p just passes back the responses to the query through QS and then through q . The Query Planner uses various Information Directories (ID_i) to understand what kinds of information are available and how to access them. Information Stores (IS_n) store, manage and access discrete bodies of information. The processes used by QP are

labeled r and s , representing the transactions that seek and retrieve relevant information needed to answer the query.

Simple model of Theory 1




Figure 1. A value chain of processing entities PE_i producing products v as a result of specifying queries and planning and executing those queries through information directories to various information sources.

Simple model of Theory 2




Figure 2. A value chain of processing entities PE_i producing products v as a result of specifying and monitoring COIs and then reacting adaptively to alerts.

The second model is very similar, and it too focuses on the same Processing Entities PE_1, \dots, PE_k that add value by accessing various information sources IS_1, \dots, IS_n to produce valued products labeled v . In this model, however, VIRT processes are at work, enabling each PE to inform the system about the COIs the system should continuously monitor. Each PE conveys its needs through interaction with a Condition Specifier (CS). The function c on the link between the PE and CS represents the transaction that yields information products consistent with PE's specification. So for example, we can assume c gives a minimal amount of information at low cost, because the PE specifies precisely what type of events it must be concerned with.

The rest of the process works roughly as follows. Once a condition is specified, the CS conveys it to the Condition Monitor (CM) through w , and CM takes responsibility for monitoring it. The transaction w just passes back any new events matching the condition through CS and then through c . The Condition Monitor uses various Information Directories (ID_j) to understand what kinds of information are available and how to access them. Information Stores (IS_n) store, manage and access discrete bodies of information. The processes used by CM are labeled r and s , representing the transactions that seek and retrieve relevant information.

Let's compare the two models. While Condition Monitoring differs a bit internally from Query Planning, the two functions use information directories and information sources in nearly identical ways. Thus the real differences between Models 1 and 2 are in the efficiency of Model 1's u , q , and p in comparison to Model 2's t , c , and w . In fact, we can simplify the analysis by viewing the valued products in each case as the output of composing the corresponding functions, so that:

$$\begin{aligned} \text{In Model 1,} \quad & v = u \circ q \circ p (K), \\ \text{and in Model 2,} \quad & v = t \circ c \circ w (K), \end{aligned}$$

where K denotes all available information and meta-data .

We read these informal equations to say, in the case of the first model for example, the value produced (v) equals what u extracts from what q extracts from what p extracts from all available information and meta-data. Equivalently, p finds relevant information in K and passes it to the next process step; this step applies q to find and pass along relevant information; finally, the process step u finds and identifies the valued information v .

For our analysis, we're assuming the valued products v are the same under the two models. These would correspond to identified threats requiring a helicopter pilot to divert from planned course, for example, or other "needles in the haystack" that operators would find worthy of selecting to act upon. In such a context, then, the two models produce the same end result, but Model 1 forces the Processing Entity to consider through process u vastly more inputs that result from query q . In Model 2, on the other hand, the Processing Entity has defined a precise COI that would materially affect expected outcomes. This means that the Processing Entity can mostly pass through results of c , eliminating any need for t to perform filtering and prioritization.

Let's consider a quantitative example. "Threats" that might affect a helicopter pilot can be natural or man-made in origin. The natural category would include high terrain, poor visibility, excessive winds, thunderstorms, or icing. Of course, these only affect the pilot if they intersect the helicopter's route of flight. Man-made threats include ground-based anti-aircraft weapons, fixed or mobile surveillance assets, and enemy aircraft. These may pose a threat if the helicopter's route intersects the volume of space these systems can observe or reach. When we speak of "intersecting," we mean the threat occupies the same space at the same time as the helicopter. So, to compare our two models, let's look at a helicopter pilot who's concerned about these possible threats intersecting the planned route of flight.

The terrain phenomena are relatively static, so there's little value in considering terrain data repeatedly, unless the pilot changes the route of flight some time after planning it. In that case, terrain previously considered benign may become a threat. On the other hand, weather changes constantly so weather data are worth considering continuously. Similarly, enemy assets may move around and allied intelligence may also improve its estimate of their positions and capabilities. For this reason, information about enemy capabilities deserves continuous consideration.

So how do the two alternative models address these needs? Each needs to look for *relevant* information and determine if it's *significant*. Information about the different phenomena is represented in various ways, but a good simplification is that all phenomena are modeled as values of appropriate variables stored in some gridded geospatially indexed array. For example, *winds aloft* are reported in terms of direction

and speed at each latitude-longitude grid coordinate, at each of several corresponding altitudes (such as 3000', 6000', 9000', etc.) The coarseness of the grid mesh differs for different types of variables in different types of systems. Usually, the geospatial data in one data array represent the expected values of the corresponding variable at a particular time. Data values that are forecast for different points in the future are stored in different arrays, each corresponding to one forecast time. When users need values that are intermediate between grid points or time points, they normally interpolate between adjacent values of the variable of interest.

In a small square-shaped theater of operations that might measure 200 km on a side, let's consider how much data is available. Let's assume that all grid meshes are 1 km, so that we have 40,000 grid points for each variable corresponding to each pair of latitude and longitude grid coordinates. We'll assume that in the vertical dimension we have data for every 500 m, from 0 (sea level) up to 6 km altitude, for a total of 13 altitude coordinates. Thus, for each variable of interest, for each moment in time, we have about 500K data values. Let's consider missions that last 4.5 hrs for which we have forecast values for each 30 min, so initially we have 10 distinct time coordinates. This means that initially, our data universe is 5M values for each variable of interest. We'll assume that we want to monitor just 10 variables in total, so that means 50 M values form our base of potentially relevant variable values. Normally, data are updated, based on new information and estimates, at least twice per hour, which means that every thirty minutes all of the data might potentially change. Because changes occur mostly asynchronously, the best strategy is to revisit the data of interest periodically, to be able to notice and respond quickly to important changes. Let's assume that our pilot decides to reexamine data every 10 minutes throughout a 4.5 hr mission.

Model 1 suggests, then, that the pilot should retrieve all relevant data every 10 minutes. The data of interest are all values of the 10 variables that intersect the planned route of flight, in the sense of overlapping spatial volume of capability with the helicopter, at the same time⁹¹. For simplification, we'll assume that 10% of the total data universe corresponds to the points in space and time where interactions might occur, if the variables indicate a threat such as a thunderstorm or enemy aircraft. So every 10 minutes, the pilot's queries access and retrieve through processes r 10% of 50M values, perhaps further reducing them by 90% using filters in q to exclude insignificant items from all the relevant data values retrieved. This means that about 1% of 50M values are returned to the pilot, or 500 K relevant and significant data values⁹². Usually, a tiny fraction of these will justify a change in plan. Most of these values, in fact, won't make a significant difference in expected outcome for the pilot. In addition, the pilot's own process u will be overtaxed by this volume of data queued for human processing. As a result, most of the data will be ignored, reactions will be suboptimal, and the pilot will feel continually stressed by the repeated onslaught of data deluges arriving every 10 min.

⁹¹ The planned route occupies a series of 3D points over time. The data of interest describe variable values at 3D points over time. The two sources are intersected over the 4D coordinates of space-time.

⁹² In today's best systems human processing capabilities are optimized by presenting much geospatially indexed data graphically as multi-colored maps. This allows the human perceptual system to process much data in parallel and detect interesting phenomena visually. We're ignoring special processing capabilities of particular machinery so we can focus on the more essential question of how to assure high-value information flows and represents an extremely high proportion of all communicated bits. All of the human's processing capabilities are limited, and they should not be squandered processing low-value inputs.

In contrast, Model 2 suggests that the processing system should take responsibility for monitoring COIs that would probably motivate and justify adaptive responses. These conditions correspond to changes in expected values of variables that the plan depends on for its success. So, for example, if there's no risk associated with faster than expected ground speeds that could result from favorable tail winds, there's no reason to monitor for tail winds. As another example, if fuel reserves are provided for 90 minutes beyond expected flight duration and are allowed to go as low as 45 minutes in all cases, only strong headwinds are worthy of considering. In fact, the system could compute the expected impact of headwinds on total flight time and only alert the pilot when the threshold of 45 minutes fuel reserve is in jeopardy. As another example, data on enemy capabilities that confirm previously considered information are insignificant to the pilot, because the system is looking for *new* threats.

Model 2 exploits its awareness of the pilot's prior and current knowledge, in addition to the pilot's plan, to drastically reduce the information passed to the pilot. The model allows the pilot to convey that understanding through process *c* where the pilot formulates specific COIs. Simplifying a bit, the PE asks the system to notify it of changes in expected values, *i.e.* events, which negate or make questionable the ability of the PE to perform its mission successfully. In the current example, over each 10 minute interval the pilot will probably receive zero or a very small number of alerts stating that some variable values have changed in significant ways. The pilot, in Model 2, has plenty of mental cycles available to handle these rare and important alerts. Furthermore, the low rate of data avoids stressing the pilot, and that further enhances human performance.

So Model 2 reduces the volume of information being communicated and also reduces the amount of work that the PE has to perform. This makes a PE in Model 2 much more efficient than in Model 1. In our example, Model 2 reduces the input volume to the PE by a factor of more than 100,000 (five orders of magnitude). If the PE in Model 1 had huge processing capacity, it could perhaps produce all desired valuable outputs *v*, just as well as in Model 2. Model 1 requires applying the “needle-in-the-haystack-finding” process *u* to about 500K items, every 10 minutes. This requires *u* to operate on 3M items suggested by the smart, excellent query *p* every hour. Typically, however, the PE is resource limited, because PE is a human and has limited cognitive bandwidth. Suppose 10% of the pilot's mental capacity is available for this task, and that humans can consider 10 significant variable values a minute. In an hour, the pilot could consider 600 reported relevant variable values, or just 1/50th of 1 % of all the retrieved information. Obviously, the pilot's response will reflect a rather random or arbitrary selection of the relevant information.

Model 2, however, requires that the pilot consider a small number of items, almost always well below the cognitive processing limits. Model 2 achieves this efficiency by conveying to condition monitoring agents a lot of “context” about the operator. For example, the Condition Monitor can take a set of assumed routes, way points, and required meteorological conditions for each pilot and continually compare and contrast the forecast weather with the required conditions. Only when weather degrades relative to a particular pilot's requirements for a specific route at a specific time and place would the CM need to transmit data back to the PE. In contrast, a system built around Theory 1 would require the pilot to ask for weather information periodically, to determine which returned values had changed significantly and, finally, to calculate which had degraded to

a point where the expectable mission outcome would no longer be acceptable.⁹³ Theory 2 suggests that all of that work should be avoided, and Model 2 shows that it can be avoided at low cost by delegating context-aware condition monitoring, as performed by CM.

Are there any situations where Model 1 is as efficient as Model 2 or even better? Yes there are. Model 2's superiority depends upon an ability to delegate to an agent the responsibility for monitoring all key assumptions expressed in terms of COIs. This requires an understanding of how outcomes depend on various variables and conditions. It also depends on capabilities to monitor those variables and compute those conditions. Furthermore, the monitoring system must be informed when plans and underlying assumptions are changed. All of this works best in contexts where plans are produced and maintained in digital, symbolic systems, and where planning processes make explicit rationales supporting decisions and choices. The military has many operations where these conditions apply, as in most tactical planning, maneuver, and logistics. However, even in these cases, the formalization of semantic concepts (ontologies), conditions of interest, pertinent information directories and information sources has not progressed very far.

So Model 1 seems to be a better fit for activities where people are engaged in less structured tasks, as when they are trying to become familiar with a new area, trying to build an initial understanding of a new situation, or when they are looking for ill defined patterns to emerge from massive amounts of data. In such situations, data mining algorithms and human intuition and perception can often be superior. Browsing, in the absence of some clear notion of what one is looking for, seems not to be a function where machines have much competitive capability.

Internet search engines, such as Google, have shown that Model 1 can be extremely valuable, especially for a wide variety of users who collectively have overlapping interests. By suggesting to the next person who asks the answers that previous investigators tended to value, search engines offer considerable advantages in comparison to human, manual, or unordered search. Model 1 seems to make concrete the idea that there are Processing Entities, usually human, that value shot-gun answers to general questions.

In contrast, Model 2 promises enormous increases in productivity for operations that are engaged in performing defined processes, repeatedly, where resources are limited and outcomes are important. In such cases, we may routinely need to let some potential opportunities "hit the floor," and there's great value in assuring effective prosecution of the most important objectives on time and within budget. In these contexts, we want our

⁹³ In this analysis, to make Model 1 as efficient as possible, we've assumed that the queries could be personalized and contextualized so that only relevant values of variables would be provided to the pilot. Thus, we've assumed in implementing Theory 1, the information universe would be honed to a small number of variables of interest and, further, that the values returned would be restricted to those where forecast values in space-time intersect with planned space-time coordinates for the flight. These assumptions go well beyond what actual systems offer or what is even being anticipated among Theory 1 practitioners. However, such improvements are logically independent of the particular Theory. We are incorporating these improvements by assumption in Model 1 to make any claims for Theory 2's advantages more vulnerable to rejection. The primary weather information systems for pilots, as just one example, afford pilots *none* of these beneficial improvements. By imagining Theory 1 is made as efficient as possible, we avoid any bias in an analysis which ultimately favors Theory 2.

people attending only to important issues. We do not want them spending time scanning, browsing, filtering, and prioritizing incoming queues that are overflowing with relevant but mostly insignificant information, or managing growing backlogs of unfinished tasks. To maximize their productivity, we want their input queues to be automatically prioritized continually so that the next item they process is, in fact, the one that has the highest expected value added. Model 2 makes that the routine process management approach.

Monitoring Conditions of Interest

When bandwidth is limited, we want to use it wisely. In any given system, bandwidth can measure the maximum possible rate of information flow through some component. Usually, of course, we think of communications bandwidth, a measure of the capacity to transport bits from one place to another per unit time. In other systems, we think of CPU clock rate or memory access rates, measures of the capacity to access or manipulate bits per unit time. As the preceding example illustrates, many processes depend on human thinking to succeed, but people can only consider a small number of variables or questions per unit time.

All systems that process information exhibit an upper limit on the amount of useful work they can accomplish per unit time, because at some level of workload, one of their components hits its maximum rate of throughput. We refer to the component that limits the total system at that point as “the rate-limiting component” or “the gating factor.” After we exceed the capacity of the rate-limiting component, we incur other problems and costs. If inputs continue to arrive, they must either be stored in some temporary buffer, housed permanently in a more expensive or slower facility or, as often actually happens, they simply “hit the floor.” In that case, the inputs are damaged or lost, never to be recovered.

When interacting with a dynamic environment or a quick enemy, it’s vital that our systems effectively focus on the most important inputs so they can implement adaptive responses before the environment or opponent creates the next problem requiring response. The basic idea in intelligent control is that one’s rate of considering and acting must be faster than the opponent’s. So, even when our systems have limited capacity, they must give priority to important information and assure their basic adaptive cycle is quick enough (Hayes-Roth, 2005c). As a result, many items may “hit the floor” when we act efficiently. However, we can’t afford to let chance determine which items the system ignores. Even when we have limited resources, we must attend to and respond to the most significant events. In our example above, the pilot in Model 1 couldn’t do this, because the deluge of inputs exceeded the resource capacity to recognize and act on the most significant information.

To integrate efficient supply chains, suppliers learn what their customers need and how they should supply inputs to reduce the customers’ difficulty, cost, and delay (*friction*) in employing them. In our pursuit of NCOW and its corresponding information chain integration, we want to apply the same principles. This encourages us to focus on the basic function of supplying information to “customers.” By analogy to the supply chain, each processing entity should supply information in a manner that minimizes the friction incurred by the intended recipient. In addition, because most information chains will be

rate-limited, every processing entity will need to attend first to the most important bits. If suppliers can prioritize information for them, processing entities can spend their limited resources more effectively, assuring they process important items before unimportant ones. In this way, the entire system can maximize the resources it expends processing the most vital bits.

To support this basic objective, we want the system to assure that processing entities receive prioritized information in a form that simplifies the recipient's task of understanding it and reacting to it adaptively. This is what Model 2 does by allowing each processing entity to specify its own conditions of interest. Each COI corresponds to a potential event that undercuts the expectation of a successful process or mission. One broad class of COIs logically corresponds to the negation of a prerequisite. For example, *sufficient fuel* is a prerequisite for flight. Each planned route considers initial fuel, fuel consumption per hour, and total hours of flight, for example. When people consider, analyze and then select a plan, they assess fuel consumption and *justify* the plan by their calculations showing the aircraft will have *sufficient fuel* to execute the plan. The corresponding most general COI is "expected fuel consumption exceeds amount available (less required reserves)." Another general class of COI corresponds to the failure of an assumed salutary condition. An example of this might be "absence of enemy threat along route of flight." While plans may not absolutely require such a condition, it's obviously desirable. In fact, in considering two alternative routes, the planner may have chosen this specific one precisely because of this very belief. These kinds of COIs reflect the importance of recognizing and adapting to events that increase the risks of failure.

A very different category of COI addresses the importance, in some operations, of adapting to increased opportunities, where we see risks as "opportunity costs," rather than chances of failure *per se*. Thus, when forecast thunderstorms dissipate, opportunities arise for shorter, faster, and safer routes.

Any NCOW system should enable processing entities to get the high-value information they need in a form that makes it easy and quick for them to react to it. This can be done by enabling each PE to specify its own COIs, reflecting its own knowledge, expected plans, and concerns about various kinds of risks. The ability to delegate and monitor COIs seems critically important. Systems intended to foster information superiority must provide this essential capability to the human operators and automated processing entities. This is probably **the single greatest way to use computing power to increase process efficiency** and to assure that limited human bandwidth can focus on important information.

Using COI monitoring this way supports the objective of every processing entity receiving *valuable information at the right time* (VIRT). Networks that embrace this process design approach will manage the flow of bits to maximize the value of those bits. We term that approach to network management *flow by value* (Hayes-Roth, 2004b, 2005a).

To enable processing entities to specify their COIs new language systems will be required, and these will constitute improved query languages. Such new language systems will require three principal components:

1. **Domain-specific ontologies.** The concepts, variables, value intervals, and similar attributes required to describe the information needed for effective missions collectively define the domain ontology. As an example, an ontology for the helicopter missions would include concepts such as *turbulence*, *icing*, *anti-aircraft missile*, *vehicle track*, *ground speed*, and *fuel consumption*.
2. **Domain-specific expressions and conditions.** Expressions specify how to calculate quantities useful in determining whether a significant value exists, as well as particular values signifying important events, *i.e.*, instances of conditions of interest. Examples would include expected *fuel exhaustion*, expected *airframe icing*, or expected *detection* by a capable anti-aircraft enemy system.
3. **Condition monitors.** Monitors are programs that take responsibility for computing conditions and alerting the interested processing entities that specified them. These monitors will access relevant data, compute the expressions associated with the conditions, and determine when the resulting values warrant notifying the interested entity or operator. In systems with extensive amounts of information, this function will be a heavy consumer of machine computing cycles. (Hayes-Roth & Brown, 2005, specifies an architecture for one such capability.)

Specialized tools to help construct and continually improve these three kinds of components will prove extremely valuable. We might consider such tools three additional types of components, making a total of six important types of capabilities required to make COI monitoring routinely available within NCOW.

Beyond the technology, organizations will need to adapt their processes as well. In the military, tactics and procedures should evolve to make COI monitoring a routine process suitable for automation. The military already has a related concept termed the Commander's Critical Information Requirements (CCIRs). Loosely speaking, these define conditions whose occurrence justifies "waking the Commander up." In today's practice, Commanders delegate CCIRs to human staff officers, and so these COIs are expressed informally in natural language.

In order to make every processing entity efficient, we look forward to the development of excellent formal systems that enable people to state their COIs precisely and simply in terms of familiar domain ontologies. In most cases, the humans should use interactive language tools that allow the computer to construct unambiguous formal interpretations while, at the same time, simplifying human performance requirements. Once formulated, the COIs would be routinely delegated to computing machines that would monitor them reliably.

We might expect these COIs to be termed something like *Dynamic Operator Information Requirements* (DOIRs). It seems straightforward to adopt a practice of explicitly defining DOIRs as a part of the planning process. DOIRs would be expressed in terms of the ontology and expressions appropriate to each domain of operation. Because all organizations spend most of their time repeatedly performing standard processes, most organizations should create a catalog of frequently used DOIRs. This would make it quick and easy to specify DOIRs for today's mission, which often would be just a slight parametric variation of items in the organization's catalog.

Discussion

I've shared earlier drafts of this paper with a number of colleagues throughout government, industry and academia. They have raised a number of points that I'd like to mention and consider here. There are four specific points that we'll consider in turn, and these are: (1) the best architecture will need to support both smart push **and** smart pull; (2) *push* and *pull* may be misleading or harmful terms; (3) people shouldn't depend on automation to provide the information they need; and (4) smart information flows require better semantics than our systems possess.

Point 1: The best architecture will need to support both smart push **and** smart pull. I've tried to emphasize to this point in the paper the significant advantage of smart push over smart pull. The earlier example showed an advantage of more than five orders of magnitude in terms of reducing information to what's essential. This is the single greatest quantitative advantage in efficiency that I've encountered in my IT career. This advantage isn't yet widely understood, and many people have a poorly informed negative bias toward smart push, mostly stemming from early Internet products that deluged enterprises frequently pushing graphical data to users. Moreover, most of the GIG/NCES and DoD policy documents emphasize access so users can seek and pull relevant data. Collectively these conditions make it important to argue against the idea that smart pull is the solution. Not only is it not the solution, it's not even a tolerable approach for many operational contexts.

However, as I suggested earlier, there are many contexts where smart push is not the answer either. Many of these contexts involve intelligence gathering or other open-ended investigations. We can be confident there are countless tasks where operators would benefit from an ability to seek and quickly find pertinent information. After all, the most useful function of the Internet today would seem to be search, which is a kind of pull. So, a broad, effective, general IT architecture will need to support both push and pull. My intent is to put smart push on the agenda and to move it to a high-priority position that reflects its enormous advantages in many operational contexts.

Point 2: Push and pull may be misleading or harmful terms. Readers have made me aware that these terms have been used in a variety of ways, ranging from technical to marketing communications, with some attendant over-loading and confusion. Hopefully, readers of this paper understand that I've been focusing on the flow of information to recipients, distinguishing *pull* flows that require recipients to periodically look for relevant data from *push* flows that cause relevant data to arrive at the recipient's inbox. Some readers point out that these functions are implemented in various ways, depending on the nature of the distributed infrastructure, and sometimes the implementations share common elements. Many of the elements of the end-to-end flows, as shown in the two process models, are similar, of course.

Nevertheless, I think it's instructive to look at the process models at a high level and consider which method provides more precise information to the operator and which method requires less work by the operator. It should be clear that two things are required to maximize the ratio of value to effort: (a) the operator's dynamic context must be used to determine precisely how current estimates differ from expectation, because those differences constitute the candidate bits of value; and (b) the system must maintain

dynamic context and past data to detect changes of interest. In addition, additional benefits accrue if the system also can project future states and detect deviations between prior expectations and now-predicted situations. Because information becomes valuable by improving an operator's expected outcome, operators will most value bits that correctly alert them to the need to change course to avoid a predictable problem or exploit an emerging opportunity.

In short, there shouldn't be any confusion about the terms, as used in this paper. Smart pull means getting the most relevant information you can when you search for it. Smart push means letting others know your context and relying on them to alert you when they determine you are at risk.

Point 3: People shouldn't depend on automation to provide the information they need.

Automation is often considered dangerous or harmful, because people may become over-reliant upon it. When automated systems fail or when novel problems arise beyond the system's boundaries, we need people to take over and provide agility and robustness. In this paper, for example, a reader might worry that our operators will become over-dependent on systems to tell them what they should know. One consequence, the argument goes, is that we weaken our operators as a result. Smart push lulls them into a kind of false level of comfort and excessive dependency.

I don't think there's anything specific to IT or smart push in such concerns. All automation ultimately creates dependencies and, in fact, people lose the skills previously required to do those functions used prior to automation. We don't have many skillful slide-rule operators today, as handheld calculators and spreadsheet software have made it possible for everybody to do complex mathematical computations at the push of a button. Few people can write database queries or Boolean information retrieval queries, yet everybody can seek and find more data, more easily, with Internet search engines. Pilots of complex aircraft regularly rely on GPS for navigation, autopilots for flying, and flight management systems to determine routing and fuel management decisions. In short, technology makes us more productive precisely because it eliminates tasks that can and often do become obsolete.

With respect to the specific question of whether smart push can provide productivity advantages, there's no doubt. Will this necessarily mean operators lose the ability to find relevant information when they want it? Will it mean that operators will no longer be effective when their systems break or are unable to address new contexts? We are certainly a very long way from having to address those questions, because we have scarcely begun to address the question of how we'd give operators just the high-value information they'd want in various dynamic contexts. At this low baseline level on the productivity scale, with little technology assistance, we know that operators can't get the information they need easily and quickly by any means. It seems obvious then that the preponderance of risk is associated with poor performance due to poor information. We should look forward to an eventual problem of having extremely high performance due to automation that provides precise, high-value information. Before operators become overly dependent and vulnerable in that eventual state of information superiority, we can introduce back up measures to increase robustness and maintain essential human skills.

Point 4: Smart information flows require better semantics than our systems possess. The examples in this paper show how smart push or pull can find relevant information and

compute conditions of interest. To do that, operators define COIs using expressions that computers can evaluate. The expressions or queries use a vocabulary of entities or variables, attributes, and values, and often these are indexed with spatial and temporal coordinates. Finally, these values might be indexed by additional aspects corresponding to different perspectives, hypothetical cases, or other context-indicating factors. A *schema* for all these data defines what kinds of bits are stored and how they're organized.

To match data to COIs, we must recognize the correspondence between the operators' concerns and the type of information modeled in the schema. In short, we need to have alignment in meaning between operators and the people who supply bits. Semantics is the term we use to refer to such meaning. As any system designer or data engineer knows, getting agreement on meaning is tough work, and the requirements continue to evolve as missions and contexts change.

So smart information flows will depend upon our machines incorporating semantic models suitable for the tasks we apply them to. The semantic models will underlie the tools users employ to express their COIs and to view the alerts and events the system pushes to them. Information will need to be registered against appropriate coordinate systems and other reference frameworks for the semantics to be meaningful and correct. Furthermore, information might be supplied using one semantic framework, compared to information in another semantic framework, and ultimately presented to operators in a third, context-sensitive framework. All of this suggests that semantics will be important, and a continuing area of learning and improvement.

There are many efforts underway throughout industry and government to create improved semantic frameworks, including ontologies and tools for comparing and translating between different category systems. I don't expect this work to move especially quickly or reach completion in the foreseeable future. On the other hand, we don't need complete solutions to begin to deliver great value to operators. We only need enough semantic foundations to understand COIs in specific operational contexts and correctly monitor evolving data relevant to those COIs. In this way, we can deliver value to operators "one thread at-a-time," by implementing the semantics-enabled processes across just those information sources needed to monitor those COIs. If we plan for continuous improvement and evolution of the semantic frameworks and semantics-supported processes, we should be able to achieve great productivity gains for the indefinite future.

Conclusions

When we don't really know what information we need or how it will change our behavior, we have few means of achieving great efficiency. In such a case we have little choice but to browse all potentially relevant information sources in the hope of noticing something interesting and, even better, something significant with respect to our situation or planned actions. This type of information system, consistent with Model 1, places a great burden on processing entities. They must seek information broadly, they must have broad access, and they must acquire large volumes of generally relevant information. Then they must process it, interpret it, look for significance, filter and prioritize it. In a network-centric operation, we should expect such processing entities to have long latencies, to induce bottlenecks, and to exhibit large backlogs of interim products

awaiting further processing. Some functions probably must be based on Model 1, such as some general intelligence operations, but we should expect efficient organizations to have few such processes.

On the other hand, we should expect mature, efficient organizations to have well integrated value chains, where each processing entity allocates most of its time to near immediate adaptive responses to *exceptions*, events that don't jibe with expectations and that jeopardize desired outcomes still in process. Efficient systems strive to achieve the best possible results by using reliable processes, each step adding value in predictable ways, and all operating within an envelope of *nominal* conditions. These systems spend most of their resources adding value by building in predictability through methods that minimize variance, reduce friction, and prevent the off-nominal unexpected. In the rare cases where conditions go off-nominal and the expected consequences are negative, the systems quickly notice them, because they are actively monitoring for just those kinds of expectable problematic conditions. In these systems, significant events occur infrequently. Further, the processing entities have mostly short queues with no backlogs, meaning they have capacity to deal with urgent events. As a result, these systems can give priority immediately to understanding a surprising event, contemplating its negative impact, and considering adaptive responses. Model 2, in short, allocates scarce attentional and problem-solving resources, often dependent on humans, in a near optimal way to recognizing and responding to vital and urgent events.

The theory of NCOW, roughly, predicts that we should be able to accomplish better results by distributing information quickly to everyone who could benefit from it. That seems like a good idea. However, Theory 1 and Theory 2 suggest two complementary strategies for achieving those improved outcomes. This paper shows that most advances in information technology, ranging from increased processing power to formal semantics, advanced query processing and knowledge-based inference, contribute significantly greater productivity gains in the context of integrated supply chains operated according to Theory 2. For this reason, organizations should give priority to designing and implementing processes as Theory 2 suggests.

References

- Ackerman, R. K. (2003). Network centrality begins at home. *SIGNAL*, August.
- Alberts, D., Gartska, J., *et al.* (2000). *Network Centric Warfare: Developing and Leveraging Information Superiority*" (2nd ed.). OSD/ASD (C3I), CCRP, http://www.defenselink.mil/nii/NCW/ncw_0801.pdf.
- Chow, Y.-W., F. A. Hayes-Roth, *et al.* (2000). Automatic retrieval of changed files by a network software agent. US Patent & Trademarks Office, #6,029,175.
- Hayes-Roth, F. (2003). Architecture, Interoperability, and Information Superiority. NPS Research Paper. Monterey, CA.
- Hayes-Roth, F. (2004a). Next-generation battlespace awareness. *Adaptive Information: Improving business through semantic interoperability, grid computing, and enterprise integration*. J. T. Pollock and R. Hodgson. Hoboken, NJ, Wiley-Interscience: 341-342.

- Hayes-Roth, F. (2004b). Model-based communication networks for improved collaboration and decision-making. Keynote Address. 12th International Conference on Telecommunication Systems - Modeling and Analysis, NPS, Monterey, DON CIO.
- Hayes-Roth, F. (2005a). Model-based Communication Networks and VIRT: Filtering Information by Value to Improve Collaborative Decision-Making. 10th International Command and Control Research and Technology Symposium: The Future of C2, McLean, VA, US Department of Defense, Command and Control Research Program (CCRP).
- Hayes-Roth, F. (2005b). Towards a rich semantic model of *Track*: Essential Foundation for Information Sharing. NPS Research Paper. Monterey, CA.
- Hayes-Roth, F. (2005c). *Hyper-Beings: How Intelligent Organizations Attain Supremacy through Information Superiority*. (Pre-publication draft.)
- Hayes-Roth, F., and D. Amor (2003). *Radical Simplicity: Transforming Computers into Me-Centric Appliances*. New York: Prentice-Hall.
- Hayes-Roth, F. and G. Brown (2005). VIRT Technical Architecture specification. W2COG Technical Working Group Document.
- Hayes-Roth, F., J. E. Davidson, *et al.* (1991). Frameworks for developing intelligent systems. *IEEE Expert* 6(3): 30-40.
- Lark, J. S., L. D. Erman, *et al.* (1990). Concepts, methods, and languages for building timely intelligent systems. *Real-Time Systems* 2(1/2): 127-148.
- National Commission on Terrorist Attacks (2004). *The 9/11 Commission Report: Final Report of the National Commission on Terrorist Attacks Upon the United States*. New York: Norton.
- Oros, Carl (2005). Helicopter Information Awareness Module (I-AM): An example of a Model-Based Communication Network (MCN) Architecture. 10th International Command and Control Research and Technology Symposium: The Future of C2, McLean, VA, US Department of Defense, Command and Control Research Program (CCRP).
- Stenbit, J. P. (2004). Statement before the Committee on Armed Services, U.S. House of Representatives, Subcommittee on Terrorism, Unconventional Threats and Capabilities. February 11.
<http://www.house.gov/hasc/openingstatementsandpressreleases/108thcongress/04-02-11stenbit.html>
- US DOD/NII TRM (Technical Reference Model) & Global Information Grid Architectures. <http://trm.disa.mil/> and <https://disain.disa.mil/ncow.html>
- Wolfowitz, P. (2004). Data Sharing in a Net-Centric Department of Defense. Department of Defense, Directive 8320.2 (December 2004), ASD (NII)/DoD CIO,

7. Valuable Information at the Right Time (VIRT) Team Mission Statement

The “Mission Statement” re VIRT

(v0.2)

Rick Hayes-Roth

Revised Nov. 23, 2005

The Problem

Information management has mostly focused till now on the question of what shall I make of the most recent snapshot of the state of my business.

On the other hand, superb competitors out-think their competition, demonstrating a superior awareness of what’s happening. They adapt to problems before the problems actually occur and thereby avoid them. They seize opportunities before their competitors perceive them. The intelligent competitor acts in a predicted future state to bring about superior results.

In the words of professional hockey’s greatest scorer, “I skate to where the puck is going to be.”

Our first problem is that current DBMSs don’t support envisioning the future and adapting to it.

The second problem is that all organizations have limited resources, especially limited time of key decision-makers. As the extent of operations increases, as the number and variety of digital information sources continue to increase, and as we look at more past and possible future states to choose winning actions, the people who have to make decisions face unmanageable volumes of information. Much organizational activity is goal-directed, plan-driven, and process-based. Productivity is directly related to how efficient each knowledge worker is in those processes. Just as lean manufacturing, JIT, and supply chain integration produced huge productivity increases by optimally scheduling the flow of materials (molecules), we need a similar transformation in the flow of information (bits).

Our second problem is to optimize the flow of bits in operational networks: how do we deliver valuable information at the right time (VIRT) to each recipient?

Our third problem is that many of the things we want to think about, monitor, and adapt have dynamic physical dimensions, so that spatio-temporal reasoning must be easy and efficient. In addition, when we think about the future, we’re often confronted with multiple alternatives as when we think about best cases or worst cases or when we must hedge against uncertain events. When we only needed to record snap-shot data, we avoided and masked the complexity of most realistic processes. If we want to use enhanced DBMS technology to represent the past, present and future states for decision-makers, we will need to support richer models of dynamic and intentional entities that operate and interact in four dimensions (space and time).

Our third problem is to enhance DMBS technology with richer models of dynamic and intentional entities that interact with each other and environmental entities. We will need to compute inferences, queries, and expressions efficiently over extremely large datasets.

Summing up: We aim to improve results by enabling each organization to adapt intelligently to problems and opportunities: by anticipating them, focusing on them, and responding smartly. We need to extend DBMS capabilities to project and reason about significant conditions of interest, especially in the future. We need systems that assure the bits decision-makers consider valuable get to them promptly and get their attention.

The Opportunity

The US Department of Defense has embarked on its own version of a next-generation Internet-enabled service-oriented architecture for the extended enterprise. The overall infrastructure is called the Global Information Grid (GIG) and the planned services are called Network-Centric Enterprise Services (NCES). This is motivated by a transformational vision of network-centric operations in which agile virtual organizations employ the GIG/NCES to put together effective processes and exploit information superiority to achieve unsurpassed competitive advantage. The total outlay for this program is likely to exceed \$10B over the next 5-10 years. Many other nations are cooperating with the US, and two consortia (NCOIC and W2COG) have sprung up to invest in and participate in these developments. Oracle participates in both.

Our colleagues working through the World-Wide Consortium for the Grid (W2COG), especially some professors at the Naval Postgraduate School in Monterey, have identified a technology agenda that would meet the needs outlined in the previous section. They have created an ecosystem that can bring together new technology, important customers, and policy makers. In addition, Prof. Rick Hayes-Roth teaches a cadre of fast-track military officers in a “capstone” masters degree course the advantages of VIRT technology as a means of achieving the DoD vision of information superiority. He’s working closely with Oracle to define our proposed roadmap. Here are some of the points he has made to us and to the leaders in DoD:

- The DoD vision cannot be accomplished without extending traditional DBMS technology to support VIRT services
- Military operators cannot afford the time or effort to sift through increasing amounts of “relevant” but *insignificant* information. In their time-stressed critical missions, they must rely on computers to do that sifting. The viability of GIG/NCES depends on implementing personalized, context-sensitive VIRT services that deliver high-value information promptly, while simultaneously filtering out low-value information.
- VIRT can reduce the amount of information that people need to process by orders of magnitude. One documented example shows an advantage of 5 orders of magnitude!

Several emerging technologies are combining to make this the right time to undertake developing VIRT services. These include:

- Semantic modeling and ontology tools (XML, XML Schemas and DTDs, RDF, OWL, Protégé, Microsoft InfoPath, and others)
- Standard off-the-shelf ontologies created by communities of interest (e.g., C2IEDM, OpenCyc, SUMO, OWL-Time, etc.)
- Symbolic reasoning, rule processing, dynamic projection, spatial and temporal computations
- Agents and service-oriented computing
- Flexible and powerful expression manipulation (making it possible to treat expressions as data)
- Low cost, high volume storage
- Low cost, high power multiprocessors, grids, and clusters
- Ubiquitous communication and mobile wireless devices

It's our belief that we can capture and exploit these emerging technologies to provide a powerful solution to the problems faced by DoD and other competitive organizations seeking to improve the speed and quality of their adaptive behaviors. Specifically, we plan to develop an extension to our DBMS products that has the following principal functions:

1. Enables an organization to express the critical assumptions that underlie its plans, so the system can monitor these. In fact, each operator and decision-maker's role in the plan is used to personalize the definition of "conditions of interest" (COI) that would significantly impact him or her.
2. Uses models of the environment and dynamic entities to describe past, present and probable future states of the world, as needed to monitor the COIs.
3. Represents COIs in user-friendly expressions and vocabulary, supported by domain-specific off-the-shelf ontologies.
4. Utilizes advanced indexing, queuing, and evaluation techniques to evaluate continuous queries for the COIs, making it possible to guarantee timely detection of COI events (changes that occur causing COIs to become true).
5. Delivers valuable information at the right time (VIRT) to each user, thereby reducing the user's information processing load by orders of magnitude and reducing communication loads proportionately.

Objectives & Goals

Our objectives are to develop and apply increasing levels of product capabilities in an iterative way, spiraling through the following dimensions as our products attain ever greater capacity, competence and performance:

- Dynamic entity models, ranging from simple vehicles to complex coordinated and distributed actors
- Domain ontologies, ranging from military and business objects to environmental features and processes
- Plans, including entity models, constraints, and assumed conditions
- Conditions of Interest, from violations of constraints and assumed conditions to achievable future states
- World Model state materialization, from continuous past to near-present to longer-term future predictions

- Spatiotemporal modeling, from 2D to 4D, from uniform grids to adaptive gridding methods that optimize storage and performance
- Expression and vocabulary tools, from SQL99 to semantic web tools (e.g., Protégé) to domain-specific semantic editors constructed atop rich semantic models and ontologies
- Implement VIRT in lighthouse applications to demonstrate the advantages in terms of faster, better decisions resulting from high-value information flows, with vastly reduced communication and processing loads on people.

Benefits

In the most profound way, VIRT should usher in a whole new generation of information management, giving the DBMS a new heart that is well suited to the network-centric environment that threatens to inundate operators and decision-makers with relevant but unimportant data. To convert those data into high-value information, we need a new system that can filter out the expected but draw attention to the unexpected significant things that signal detection of new problems or new opportunities.

This means we need a technology suitable for managing the future, not merely reacting to the last snapshot of the past.

This, in turn, means we need to be able to understand and monitor important conditions as dynamic events unfold. To do that, we will enable our customers to describe the dynamic entities in their application domains and to tell us how best to predict or obtain forecasts for those entities. Our next generation database will be able to represent these time-indexed states of important entities, variables, and expressions, and this will enable us to continuously evaluate when important conditions are about to occur. By alerting our customers to those anticipated events, we enable them to adapt before the opportunity is lost.

In short, we are creating a capability for all organizations and users to leverage advancing IT to improve their ability to foresee and control events.

By assuring that our systems convey just important information to people, we are also making the exploding information environment hospitable. We are reducing the information glut by giving focus and priority to news that significantly affects the recipients. That eliminates the need for them to receive and sift through huge amounts of information, something that is increasingly impossible.

In short, our VIRT project should mark a sea change. Our current products focus on recording and reporting on the recent past so that managers looking backwards might attempt to improve the future. Our new products will directly enable managers to see the impending future, to hear when it violates their critical assumptions, to generate and evaluate potential changes in plans, and to select optimal courses of action. It's no overstatement to say that the new products will bring us into the age of proactive, intelligent enterprise management.

Technical Work Required

We must work through the technical dimensions described with a spiral approach, climbing to ever higher levels of capability on interacting dimensions, over successive spirals. As our first spiral, we aim to develop the following levels of capability:

Spiral 1 (built up in three successive increments, each 4 months optimally).

- Dynamic entity models: rich semantic tracks (entities that move over time and space in predictable ways), for tracking vehicles, people, parcels, containers, etc. Also applicable to environmental entities such as weather fronts, hurricanes, etc.
- Domain ontologies: specific entities that instantiate tracks; also basic transportation networks where tracks move.
- Plans: plans for such tracks, and plans for managing the transportation networks.
- Conditions of Interest: tracks deviating from acceptable behavior profiles; likely conflicts between tracks.
- World Model state materialization: assuming excellent adaptive decisions can be made whenever we have T minutes advance warning of COI events, we will project ahead 2T, materializing the World Model state using adaptive grids and lazy computations.
- Spatiotemporal modeling: We will provide both rich spatial modeling for surface of the earth and coarse modeling to 3D spaces for aviation applications, with time scales proportionate to decision cycle time T.
- Expression and vocabulary tools: We'll develop a flexible syntax for expressions, built atop a flexible meta-model for vocabulary terms, so that we can provide a flexible, easily modified stack of tools that culminates in domain-specific form-filling end-user editors for COIs and vocabulary. We will make these work with Microsoft Office (InfoPath).
- Implement VIRT in lighthouse applications (TBD)

Models, Semantics, Inference

As you know, the Semantic Web is a hot topic and is anticipated to be a driver for vast amounts of computing cycles. Most of the W3C effort on the Semantic Web is addressing broad classes of low-value information, where the technology aims to identify more relevant information without being able to identify significant information. Our focus is to add some additional depth to what is modeled so that we can understand, for each user, what is important. All current efforts to provide better decision-making tools focus on understanding what the data *means*, which is the subject of *semantics*. Semantic models range from simple agreed vocabularies and schemas to rich, rigorous logical systems with commonsense knowledge of the world (as in Cyc). We believe that one way to hit a bulls-eye on detecting and delivering important information is to detect when people's plans are at risk because their critical assumptions are **turning out to be wrong**. We believe that detecting such critical assumption violations will be highly valued by many customers. So, this guides our focus on what we must model and what kinds of computations we need to do on these models.

As suggested in the list of Spiral 1 objectives, we intend to **model** the following things:

- States of dynamic entities, from the past to the projected future

- Space and time

- Intentional beings, especially those operating or traveling in vehicles

- Physical things, especially vehicles that transport other things or carried in vehicles

- Environmental things, that interact with the intentional beings and physical things

The types of **inference** that we must support include:

- Relating observations and reports to entity models, so that we “track” things.

- Materializing the future states of dynamic entities, from most probable to merely likely.

- Determining when tracks and COI-relevant spatial envelopes intersect.

- Determining when important variables attain particular values, in some space-time region of interest.

Over time, we expect to leverage worldwide work on building reusable ontologies for other domains. We are grounding our work on problems common to many defense, government, and civilian management problems, where plans are common, where future forecasting is routine, and where we can establish significant first mover advantage.

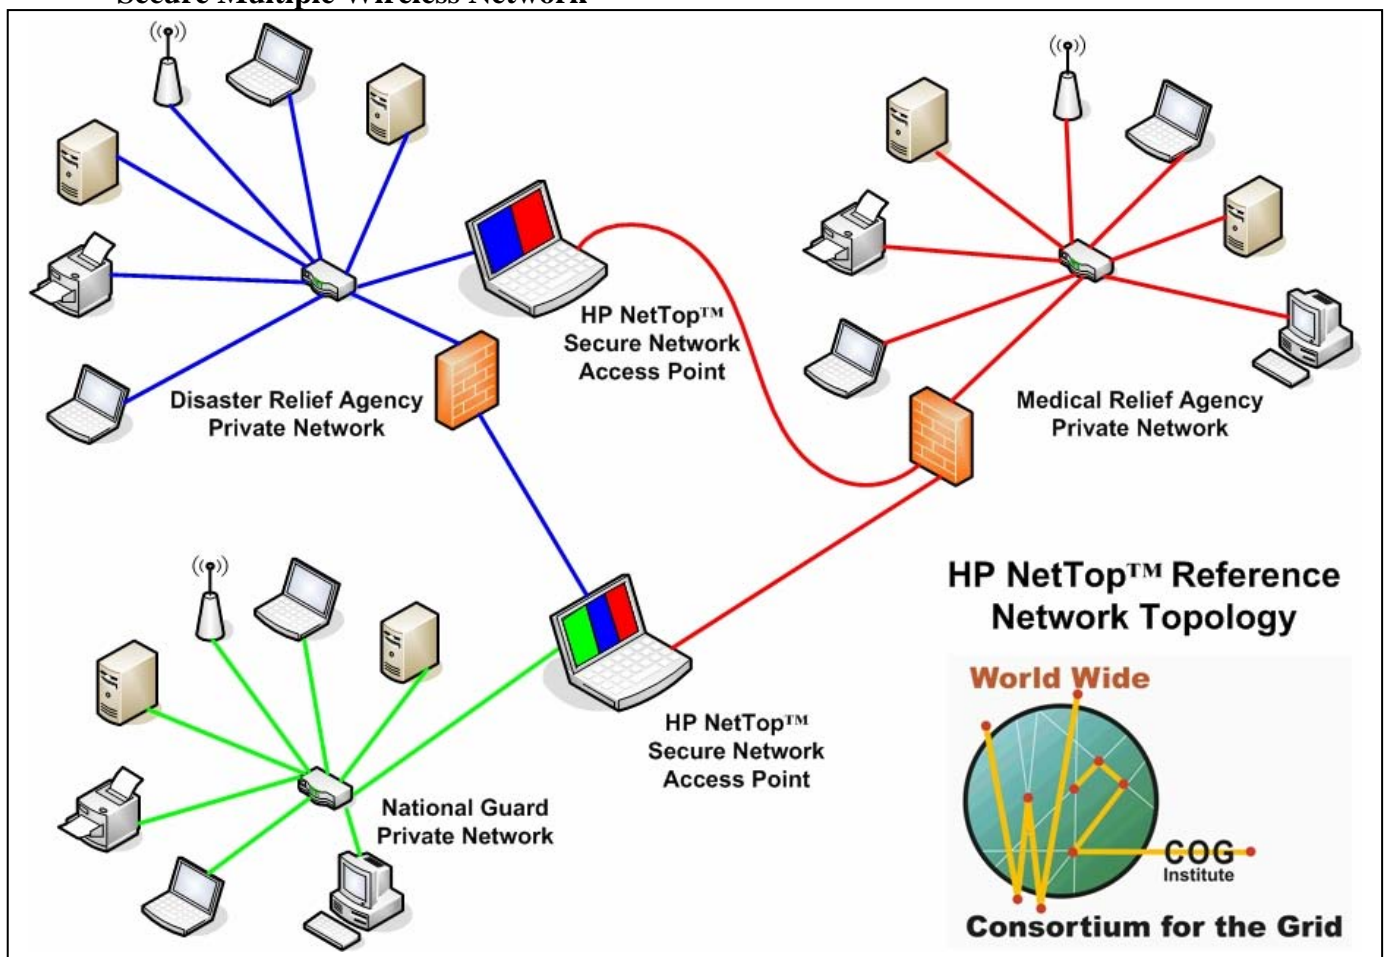
8. Reference Implementation Documentation of Humanitarian Privacy-Preserving Collaborative Network

HP NetTop™ Demonstration

Recent environmental disasters such as Hurricanes Katrina and Rita have highlighted the requirement for establishing hastily formed networks by the responding relief organizations. The first-response organizations must have the ability to establish a private Local Area Network (LAN) and secure Wireless LAN (WLAN) to enable information sharing/gathering in support of the relief mission. Some of the information gathered could be of a sensitive nature, such as Social Security Numbers and/or Medical information, etc., while other information could be of a collaborative nature, such as location, capabilities and/or supplies, etc. Communication infrastructures can become intermittent or unreliable based on the magnitude of the event. The creation of secure WLANs could allow first-responders to establish communication without relying on typical communication infrastructure.

The relief organizations must have the ability to share collaborative information enhancing the mission, while maintaining the security and integrity of the private information on the network. HP NetTop™ allows a single computer to access different network security domains. The separation can provide access to sensitive internal information and external resources to enable safe interaction between coalition and other partners. HP NetTop™ provides an environment supporting typical COTS productivity software, enhancing worker efficiency while maintaining the separation of sensitive information located on isolated networks. The HP NetTop™ architecture could be expanded to enable the addition of wireless connectivity capabilities. Secure multiple wireless networks could enhance the mobility of first-response personnel, thereby increasing the effectiveness and decreasing the time-to-delivery of critical services.

Secure Multiple Wireless Network



NetTop is a trademark of the National Security Agency.

HP NetTop™ Reference Installation


Background:

The National Security Agency (NSA) researched methods to provide cost-effective high-assurance applications using commercial-off-the-shelf (COTS) technology. NetTop™ was developed by the NSA to provide an architecture capable of executing multiple concurrent operating systems (OS) while maintaining complete isolation between the systems. Each of the concurrent OS can be attached to a different security domain; effectively creating isolated workstations with access to multiple isolated networks on a single computer.

Each of the OS is associated with a different network through a dedicated network interface card (NIC). The NetTop™ software establishes a policy-based virtual air gap by preventing data crossover between security domains. The separation and isolation provided to each OS by NetTop™ allows the operator to use COTS productivity software that is typically unavailable to high-assurance installations.

The Hewlett-Packard Company, in a licensing agreement with the NSA, offers HP NetTop™ as a full-service solution to public and private enterprises. HP NetTop™ is a highly secure, layered environment of SELinux, VMware™ and customized security policies. It is backed by the HP Technology Solutions Group to provide assessment, planning, policy definition, rollout, and support tailored to your organization.

Ultimately, HP NetTop™ is simple and transparent to the end user. The user simply clicks between OS windows just as they do between application windows. The underlying NetTop™ software works in the background to maintain system isolation. Each guest OS executes in its own virtual machine (VM). The VMs are indistinguishable from a standalone workstation on the Internet or company intranet. Applications work transparently within the HP NetTop™ VMs. The VMs run independently from the host NetTop™ OS and can crash without affecting other VMs or the host OS.



Two Enclave HP NetTop™ Installation

In a typical installation like the one shown above, each VM is connected to a different network to separate and isolate the applications from the other VM. HP NetTop™ is effectively a secure software KVM switch for virtual machines. This document will identify the hardware and software requirements for duplicating a specific two-enclave reference installation. Additionally, a hardware specific instruction set will be added to supplement and highlight the pertinent vendor instructions.

References:

- HP NetTop™ 1.2 Installation and Configuration Guide
- HP NetTop™ 1.2 Users Guide

Hardware Specification:

The HP NetTop™ software operates on a variety of hardware platforms. Other models are documented in the HP NetTop™ Installation and Configuration Guide including other notebook and desktop models. The following are the hardware specifications for the reference implementation, which include a notebook computer and a PCMCIA NIC.

- Compaq nc8000 Notebook Computer (p/n PF076US#ABA)
 - Upgrade to 1GB of RAM
- One of the following PCMCIA NIC:
 - 3Com 3C574-TX Fast Etherlink PCMCIA Card with dongle (p/n: 16-0096-000 Rev:A)
 - 3Com 3C589D EtherLink III PCMCIA Card with dongle (p/n: 16-0037-000 Rev:A)
- Any USB Mouse

Software Specification:

- HP NetTop™ version 1.2 software installation CD
 - Contact: nettop@hp.com
 - Further information: www.hp.com/go/nettop
- Software license for guest operating system for each virtual machine (VM)
 - Example: (2) Microsoft® Windows® XP Professional licenses
- Software license for any additional software
 - Example: (2) Microsoft® Office Professional Edition 2003 licenses

Hardware Specific Supplemental Instructions:


Please read the HP NetTop™ 1.2 Installation and Configuration Guide (1.2 ICG) and the HP NetTop™ 1.2 Users Guide (1.2 UG) carefully. No other special hardware, databases, software or specialized testing is required for the reference implementation. Other installations may require databases or other software to provide the functionality required as the circumstances dictate. The majority of the items included below are in the 1.2 ICG but are added here for clarity.

- The PCMCIA NIC must be installed prior to installation. The topmost VM listed on the Start menu binds to *eth0*, which is the laptop Ethernet adapter.
- Once to the configuration menu, select Video Configuration and select (5) custom. Select 1400x1050 resolution, 32 bit depth, use the default sync ranges, Analog and under Dualhead select No. (The default 15 inch LCD/Laptop (4) will work as well but the screen size is more functional at the larger resolution. Other settings may be appropriate depending on the specific installation.)
- Go into the mouse configuration and set for USB rather than PS2 and start the installation.
- Note the instructions for associating the CD to the individual VM in the 1.2 UG.
- In the Troubleshooting section of the 1.2 ICG there are instructions if “installing a Windows guest OS is extremely slow”. Follow the instructions and the installation will be quicker.
- Do not be concerned with poor graphics resolution after installing the Windows OS and rebooting. Continue to follow the instructions and install the VMware Tools. The screen resolution and display should be normal after installation of the VMware Tools.
- Once the guest OS has been installed in each VM, the administration and installation of secondary software is handled as if each VM was a completely separate computer.
- Once all settings are satisfactory, continue to follow the instructions in the 1.2 ICG and make the system secure. Do not make the system secure until all functionality has been verified and no modification to the installation will be required.

9. Business Transformation Agency Enterprise Software Incubator Proposal

World Wide Consortium for the Grid (W2COG)

1895 Preston White Dr
Reston VA 20191
(202) 281 8872
aaron.budgor@w2cog.org
www.w2cog.org



January 23, 2006
via email transmission David.Scantling@DFAS.MIL

Business Transformation Agency
Attn: Mr. David Scantling
1851 S. Bell St., Suite 1002
Arlington, VA 22240

SUBJECT: W2COG Institute Proposal for Support to Business Transformation Agency

Dear Mr. Scantling:

Please accept this input for an Executive Summary and supporting slide deck describing how the W2COG Institute can support the Business Transformation Agency enhance DoD business enterprise visibility.

The principle elements of this proposal include:

- W2COG Background and Accomplishments
- Development, Execution, Validation and Value-Add of Lighthouse Pilots
- Conduct of Case Studies on Business Processes and Policies
- Justification for Series A Funding

I and the membership of the W2COG Institute look forward to becoming an integral part of your Agency and to working with you and the BTA on this exciting and important project.

Sincerely,

Aaron Budgor
Managing Director, W2COG Institute

W2COG Background and Operation Construct:

The Global Information Grid (GIG) represents a fundamental shift by the DoD in information management, communication, and assurance to meeting the timely needs of both the warfighter and the business user. Senior OSD leadership recognized that implementing this vision would require vastly higher levels of collaboration across heretofore autonomous organizations and would need leveraging of commercial technologies augmented to meet DoD's mission-critical user requirements. They also recognized that the current DoD acquisition landscape neither provides incentive to nor convenient processes to encourage cross domain collaboration. They were encouraged by successes in the e-business private sector that have found ways to adopt collaborative practices to achieve competitive advantage in the marketplace.

Accordingly, the Office of Force Transformation, ASD Networks & Information Integration, DUSD Advanced Systems and Concepts, and DARPA, collectively provided \$1.7M "angel money" to the Naval Postgraduate School to establish the World Wide Consortium for the Grid (W2COG) research initiative. The project objective is to create a self-sustained not-for-profit open consortium that applies the Internet "open" e-business process to accelerate the GIG.

The W2COG research initiative discovered that the principles of netcentric operations (NCO) can be effectively applied to engineering. That is, self-synchronized teams of vendors taking their cue from the Open Source movement can rapidly bundle their separate products to create incrementally more powerful information processing capability. This idea led to two central themes: The first is Valuable Information at the Right Time (VIRT); the power of NCO is in enhancing information utility, not moving data. The second is Value off the Shelf (VOTS); the power of netcentric engineering is in creatively re-using interoperable (i.e. off-the-shelf) components. Hence, W2COG projects demonstrate quantifiable improvements in information processing capability by bundling excellent off-the-shelf components.

The W2COG research initiative spawned the not-for-profit W2COG Institute in July 2005 to establish and maintain the infrastructure required to achieve the goals of the W2COG vision.

The tenets of the W2COG Institute are to 1) create a forum and facility to discover commercial and government best practices and solutions for network and collaboration technologies; 2) establish and maintain a readily searchable data base of government, industry, and academic experts in operational, engineering, and programmatic aspects of net centric operations required to create the solution(s) required; 3) demonstrate the utility of off-the shelf network and collaboration components, and how they can be bundled and used to rapidly satisfy NCO requirements, to include placing them on commercial or government procurement schedules; 4) produce documentation that accompanies the successfully demonstrated bundled capability, and to perform independent testing and validation of the solution ("Underwriters Laboratory" function); and 5) sponsor Research and Development projects and provide grants in areas pertinent to networking and collaboration technologies and their Independent Validation and Verification.

Achievements to date include provisioning of an IPv6 transition strategy to the Australian Defense Force; fielding a reference implementation of an unclassified, multi-level security domain that allows disparate organizations responding to a humanitarian disasters to share some information while keeping other information private; a functioning prototype architecture that applies semantic web technology to the DoD's C4 "track" construct; acquisition of a COTS/GOTS Sensor-Net for Perimeter Protection for US forces in Iraq, and deployment of Hastily Formed Network solution for Katrina relief and distance learning. Potential new starts being negotiated include:

- DISA
 - Review of GIG SOA standards
 - Assessment of DT&E and OT&E requirements on COTS procurement
 - Pilot efforts on ECM and on Applications Services

- National Guard Bureau

- Operations Center support in collaboration with NORTHCOM, DHS and state and local governments
- NASA & FAA
 - Assistance in leveraging GIG development for federal aviation applications
- Joint and Coalition NCO experimentation through JFCOM and EUCOM R&E relationship with NPS

W2COG Institute Value Proposition to Business Transformation Agency

The W2COG Institute proposes to support the BTA by organizing and managing projects to enhance DoD business enterprise visibility by delivering reference implementations in well defined increments of DOTMLPF. W2COG Institute recognizes that application of commercial techniques will necessitate speed to market; thus it recommends that each pilot be based on a maximum duration of 6-9 month refresh cycles.

The W2COG Institute proposes to develop criteria for pilot selection and governance and apply these criteria to at least one lighthouse pilot chosen to demonstrate financial and/or warfighter business enterprise visibility. The financial pilot, might, for example be a COTS application of SFIS to manage funds allocation, collection, control and disbursement. The warfighter pilot might make use of RFID technology on the Army GSM network for a logistics application in Iraq; and incorporate wired and wireless infrastructure, to include GIG-BE, Bluetooth, WIFI, WIMAX and SATCOM. The essential ingredients of such “Lighthouse Pilots” include: operators and engineers working together; quantitative productivity targets; and reference documentation to support government-approved “user’s manual”.

Finally the W2COG Institute proposes to Conduct and manage Pilot(s) and then assess their value-add to current and future Programs of Record. Pilots will deliberately explore ways to apply consumable and/or service contracts to rapidly field successfully demonstrated capability.

To accomplish this work, the W2COG Institute requests that the BTA provide Series A funding to carry on the significant work begun during the W2COG 18 month start-up phase. We estimate that to accomplish the above described tasking would require funding of \$875 K over one year. These funds would be used to provide salaries and other direct cost expenditures dedicated to W2COG Institute employees executing and providing administrative support of the above tasking; delivery of one lighthouse pilot; and providing two case studies on Business Processes and Policies. Project leadership and technical performers are comprised of experts and thought leaders from government, industry and military operators chosen from Consortium membership based on past performance and peer review. Additional expertise for any part of this project will be solicited, as required.

We require two distinct funding tracks; one will cover salary of the government project officer, and the other will support the non-government W2COG Institute. Due to the non-profit status of the Institute we recommend funding in the form of a grant to the W2COG Institute. Options to cover the government project officer include direct hire, IPA, or research grant to the Naval Postgraduate School.



World Wide



COG
Institute

Consortium for the Grid

Proposal for Support to DoD Business Transformation Agency

www.W2COG.org

10. Reference Implementation Documentation of GLG-lite On-line Environment

REFERENCE IMPLEMENTATION: *GIG-lite Netcentric Test & Evaluation, Validation & Verification Runtime Environment*

TAB A: Equipment, Data, and Interfaces

TAB B: How to Publish to a GIG-cast Channel

TAB C: Mission Thread Descriptions

Summary: The GIG-lite is a web service oriented, IP based, runtime environment designed for demonstration and evaluation of network enabling services. Geographically distributed developers can easily configure computer network-related artifacts over the Internet via standard open interfaces, and evaluate their performance against customer designed mission threads (i.e., use cases) in an environment that simulates actual conditions. The intent is to demonstrate sufficient computer network utility and security, while measurably enhancing productivity via these deployed artifacts. Artifacts will be combinations of hardware, software, tactics, techniques, processes, policies, etc. GIG-lite is designed to encourage co-evolution of technology, doctrine, and policy. Customers of computer network capabilities can use the GIG-lite as an on-line shopping service, per e-Bay and Amazon.Com models. A reference implementation of a rudimentary GIG-lite was delivered to the W2COG Working symposium 24 May 2005. It was hosted on a small suite of generic PCs connected via local IP backbone. Functionality was provided by three existing DoD C4ISR systems, Precision Air Drop System, NetTop multi-level security thin client, and Commercial Web-enabled Service Toolkit (CWEST). Demonstration scenarios were a Joint Special Operations Task Force Intelligence cell “tipper to target” problem, a humanitarian disaster air-logistics problem, and a international border crossing identity validation problem.

Configuration. Figure 1 is a logical sketch of the GIG-Lite reference implementation. TAB A lists GIG-lite equipment, interface documentation, and data types. TAB B explains procedures for publishing and subscribing to services. Mission scenarios, use cases, and desired capability enhancements are described in TAB C. All of the capabilities available within the GIG-Lite are currently operational in various US DoD C4ISR applications or operations. Systems represented are:

Precision Air Drop System (PADS). PADS uses computer generated predictions to estimate wind fields at a target area. In situ data is then collected during the operation and the wind fields are updated to provide a higher resolution estimate of the winds. This estimate is then used to compute a Computer Air Release Point (CARP) based on the payload, altitude and speed, aircraft and station.

NetTop. NetTop allows a user to display, concurrently, on the same desktop, two levels of security separation. In the reference implementation demonstration, these are treated as track data at the Top Secret SI level and tracks at the Secret coalition level. Although these are operationally separate enclaves, for purposes of ease of operation, they are shown here on the same network—however, the interfaces to the system are separate. There is no actual classified data in this GIG-lite reference implementation, only

test sets are being used. The demonstration provides Common Operational Picture COP views for track data at a simulated TS SI and Secret Rel.

Commercial Web-enabled Service Tool (CWEST). CWEST allows users to set up a GIG-cast “channels” interface for a publish/subscribe implementation that can support arbitrary payloads.

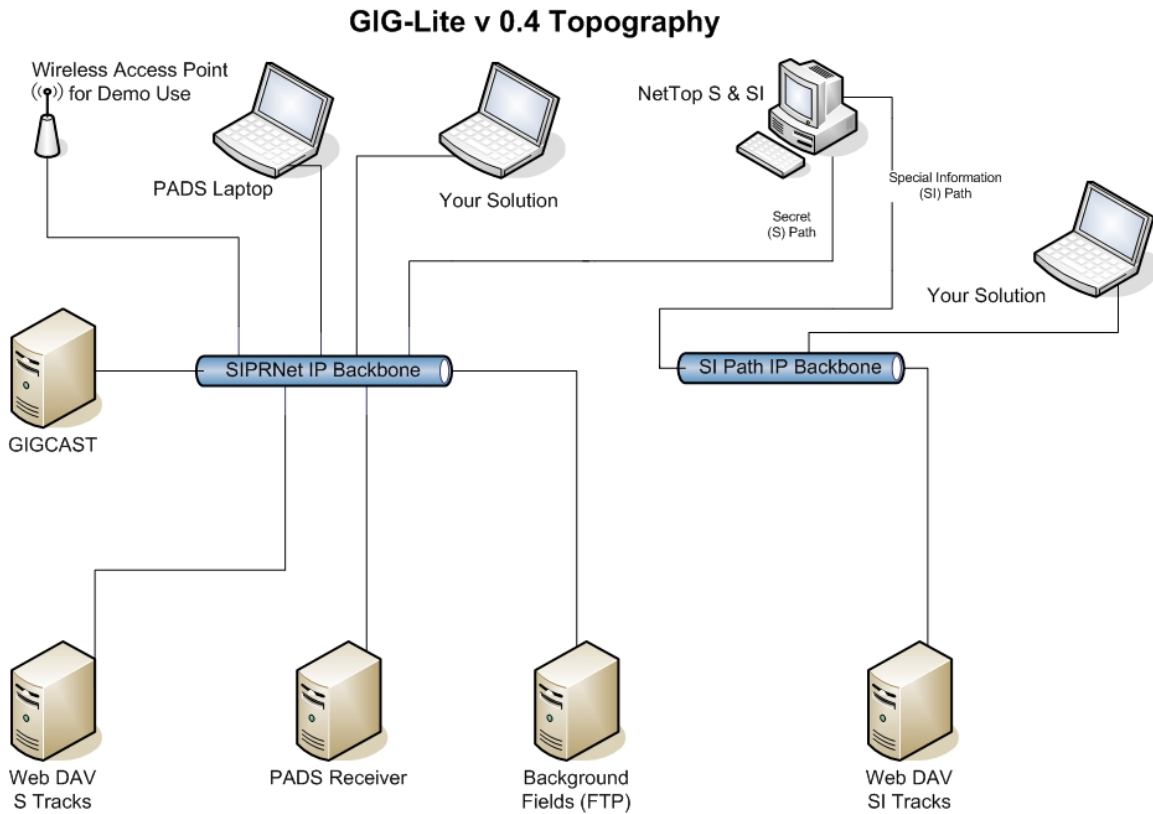


Figure 1: GIG-Lite Topography

Availability: The GIG-Lite will be deployed on the W2COG.org web site and available to members. GIG-lite will grow and improve continuously as a distributed resource.

TAB A: Equipment, Data, and Interfaces

GIG-Lite v 0.6 Interface Documentation

Host Machine Name	Equipment	Operating Systems	Mission Thread
Background Fields	Dimension T800r: 800MHz Pentium III, 256MB PC100 SDRAM, 40GB HDD	Red Hat Enterprise	PADS, Your Solution
PADS Receiver	Dimension T750r: 750MHz Pentium III, 256MB PC100 SDRAM, 40GB HDD	Red Hat Enterprise	PADS, Your Solution
Secret (S) Tracks	Dimension T800r: 800MHz Pentium III, 256MB PC100 SDRAM, 40GB HDD	Red Hat Enterprise	NetTop S Path, Your Solution
Special Intelligence (SI) Tracks	Dimension T800r: 800MHz Pentium III, 256MB PC100 SDRAM, 40GB HDD	Red Hat Enterprise	NetTop SI Path, Your Solution
GIGcast	Optiplex GX110: 933MHz Pentium III, 256MB PC100 SDRAM, 40GB HDD	Red Hat Enterprise	PADS, CWEST, Your Solution
Sun Blade	Sun Blade 100: 500MHz Ultra SPARC-Ile 512MB RAM, 15GB HDD	Solaris 10	Your Solution
GIG-Lite Client			
NetTop S Path	HP D530	Linux SE Host Win 2K (VMware)	Win 2K OS available for Your Solution
NetTop SI Path	HP D530	Linux SE Host Win 2K (VMware)	Win 2K OS available for Your Solution
PADS Laptop	Panasonic CF-29	Win XP Host/RH Linux (VMware)	N/A for Your Solution
Your Solution	TBD	TBD	TBD

GIG-Lite v 0.6 Interface Documentation

Host Machine Name	Data	Interface	Access Control	Example Files	Supported Mission Thread	IP Address
Background Fields	MM5 Gridded Binary (GRIB)	FTP	User: GIG, Password:GIGus057g1010t02g060000600		PADS, Your Solution	192.168.10.11
PADS Receiver	ECEF Formatted wind data Specific to PADS Laptop	HTTP	None	Streaming Data	PADS, Your Solution	192.168.10.12
Secret (S) Tracks	XML	Web DAV	User: GIG, Password:GIG	GCTF1!HADRData[1].xml	NetTop, Your Solution	192.168.10.10
Special Intelligence (SI) Tracks	XML	Web DAV	User: GIG, Password:GIG	4EYES!SOFsensors[1].xml	NetTop, Your Solution	192.168.2.1
GIGcast	Available Channels	Web API	None	TBD	PADS, CWEST, Your Solution	192.168.10.13
Sun Blade	TBD	TBD	TBD	TBD	Your Solution	192.168.10.14
GIG-Lite Client						
NetTop S Path	XML	None	User: GIG, Password:GIG	N/A	N/A	192.168.10.20
NetTop SI Path	XML	None	User: GIG, Password:GIG	N/A	N/A	192.168.2.20
PADS Laptop	GRIB, ECEF wind	None	None	N/A	N/A	DHCP
Your Solution	TBD	TBD	TBD	TBD	TBD	TBD

GIG-Lite v 0.6 Interface Documentation

Host Machine Name	Data	Format	References
Background Fields	MM5 Forecast Fields located in folder MM5_Forecast_28Apr05	WMO GRIB	https://mel.dmsomil/docs/grib.pdf
PADS Receiver	ECEF Formatted wind data Specific to PADS Laptop	HTTP	N/A
Secret (S) Tracks	Ship Track	XML	http://www.polexis.com/pdfs/xis_slick_11-02.pdf
Special Intelligence (SI) Tracks	Ship Track	XML	http://www.polexis.com/pdfs/xis_slick_11-02.pdf
GIGcast (Ocean Channel)	Bathymetry Data, JOTS-W Global Warnings	Bathy: XML, JOTS-W: OTH	See attached
GIGcast (Atmos Channel)	Surface Temp, Surface Pressure, 12_hr Precipitation	WMO GRIB	https://mel.dmsomil/docs/grib.pdf
GIGcast (MM5 Forecast Channel)	See included MM5 GRIB description for parameters.	WMO GRIB	https://mel.dmsomil/docs/grib.pdf
GIGcast (PADSWind)	Assimilated in-situ wind data, see PADSWind description.	ASCII Text, Tab Delimited Format	See attached
Sun Blade	TBD	TBD	TBD

Specific MM5 Parameters for PADS GRIB Files

Parameter Description	GRIB File Parameter Names from t2z MM5 Trimgrib Datasets:	GRIB ID	Level
2D Latitude	Latitude	230	1
2D Longitude	Longitude	231	1
3D Geopotential Height	Geopotential Height	7	100
3D Mixing Ratio	Mixing_ratio	53	100
3D Temperature	Temperature	11	100
3D Vertical Winds	Pressure_Vertical_velocity	39	100
3D E-W Winds (U)	u-component_of_wind	33	100
3D N-S Winds (V)	v-component_of_wind	34	100
Surface Pressure	Pressure_Vertical_velocity	1	1
Surface E-W Winds (U)	u-component_of_wind	33	7
Surface N-S Winds (V)	v-component_of_wind	34	7
Surface Temperature	Temperature	11	105
Surface Dew Point	Dew_point_temperature	17	105
MSL Pressure	Pressure_Reduced_to_MSL	2	102
Terrain Height	Model_Terrain_Height	233	1
Altimeter Setting	Altimeter_Setting	209	105
Grid Geometry: Projection Type. Two "true" latitudes for Lambert-Conformal. Center longitude (where latitudes are N-S). Number of grid points.			
Grid spacing.	GRIB header info		
Forecast times	GRIB header info		

Note: for 3D parameters, the levels are 1025, 1000, 975, 950, 925, 900, 875, 850, 800, 750, 700, 650, 600, 550, 500, 450, 400, 350, 300millibars

Wind Profile	Forecast & Dropsonde	
	6-Nov-	
Planned Date:	03	DD-MMM-YY
Planned Time:	19:50Z	HH:MMZ
Planned PI Lat:	N 33_22.160	[N/S] DD_MM.MM [Hemisphere Degrees_Minutes.Decimal Minutes] [E/W] DDD_MM.MM [Hemisphere Degrees_Minutes.Decimal Minutes]
Planned PI Lon:	W 114_16.513	
Planned PI Elev:	1360	Feet above MSL [Planned Point of Impact Terrain Elevation]
Sonde IP Lat:	N 33_22.256	[N/S] DD_MM.MM [Hemisphere Degrees_Minutes.Decimal Minutes] [E/W] DDD_MM.MM [Hemisphere Degrees_Minutes.Decimal Minutes]
Sonde IP Lon:	W 114_16.761	
Sonde IP Elev:	1354	Feet above MSL [Sonde Impact Point Terrain Elevation]

Altitude	Wind Profile		Ballistic Wind			
Feet	Direction	Speed	Direction	Speed	Lat	Lon
	Deg		Deg			
MSL	(Mag)	Knots	(Mag)	Knots		
0	999	999	999	999	999	999
1000	999	999	999	999	999	999
* 1360	297	4	297	4	N 33_22.160	W 114_16.513
2000	316	5	308	4	N 33_22.160	W 114_16.513
3000	352	3	319	4	N 33_22.160	W 114_16.513
4000	141	6	327	2	N 33_22.160	W 114_16.513
5000	129	2	8	0	N 33_22.160	W 114_16.513
6000	230	2	163	0	N 33_22.160	W 114_16.513
7000	217	5	215	1	N 33_22.160	W 114_16.513
8000	206	5	213	1	N 33_22.160	W 114_16.513
9000	204	7	210	2	N 33_22.160	W 114_16.513
# 9429	204	8	209	2	N 33_22.160	W 114_16.513
##						
10000	208	9	208	3	N 33_22.160	W 114_16.513
11000	220	14	211	3	N 33_22.160	W 114_16.513
12000	231	20	216	5	N 33_22.160	W 114_16.513
13000	238	28	222	6	N 33_22.160	W 114_16.513
14000	238	32	227	8	N 33_22.160	W 114_16.513
15000	236	34	229	10	N 33_22.160	W 114_16.513
16000	233	35	230	12	N 33_22.160	W 114_16.513
17000	230	35	230	13	N 33_22.160	W 114_16.513
18000	228	35	230	14	N 33_22.160	W 114_16.513
19000	227	37	230	16	N 33_22.160	W 114_16.513
20000	227	39	230	17	N 33_22.160	W 114_16.513
21000	226	41	229	18	N 33_22.160	W 114_16.513
22000	224	43	229	19	N 33_22.160	W 114_16.513
23000	221	45	228	20	N 33_22.160	W 114_16.513
24000	219	47	227	21	N 33_22.160	W 114_16.513
25000	217	50	227	22	N 33_22.160	W 114_16.513
26000	215	52	226	24	N 33_22.160	W 114_16.513
27000	213	54	225	25	N 33_22.160	W 114_16.513
28000	211	57	224	26	N 33_22.160	W 114_16.513
29000	211	60	223	27	N 33_22.160	W 114_16.513
30000	211	64	222	28	N 33_22.160	W 114_16.513
31000	211	66	221	29	N 33_22.160	W 114_16.513
32000	211	66	220	31	N 33_22.160	W 114_16.513

33000	211	66	219	32	N 33_22.160	W 114_16.513
34000	211	66	219	33	N 33_22.160	W 114_16.513
35000	211	66	218	34	N 33_22.160	W 114_16.513

* Planned PI Elevation

Sonde First Wind ALtitude

Aircraft Altitude at Sonde Release

Joint Operational Tactical System (JOTS) Message Format

JOTS Warnings JOTS-W File

A JOTS Warnings File is a collection of warning messages in Over-The-Horizon (OTH) format, preceded by a file header (created by a transmission protocol, typically DPSR)

```
<JOTS-file> ::= <DPSR-header>
               <OTH-message>*
```

A transmission file header <DPSR-header> created by DPSR is a set of lines. The first line contains one word "BEGIN"; the last line of the <DPSR-header> has a single word "END" (immediately followed by a newline).

An <OTH-message> is a collection of lines; each line is exactly 69-characters long (space-padded if needed) terminated with a #\newline character.

Line's

content is entirely in upper case.

```
<OTH-message> ::= <OTH-header> <body-lines>* "ENDAT" 64#<space> "\n"
<OTH-header> ::= <message-id> " " <timestamp> <space-padding>
<message-id> ::= 7-alphanum characters, starting with "METX"
<timestamp> ::= <JJ><MM><YY><GG>
```

where <JJ> are the last two digits of the year, <MM> is the UTC month of the year, <YY> is the UTC day of the month, and <GG> is the UTC hour.

TAB B:

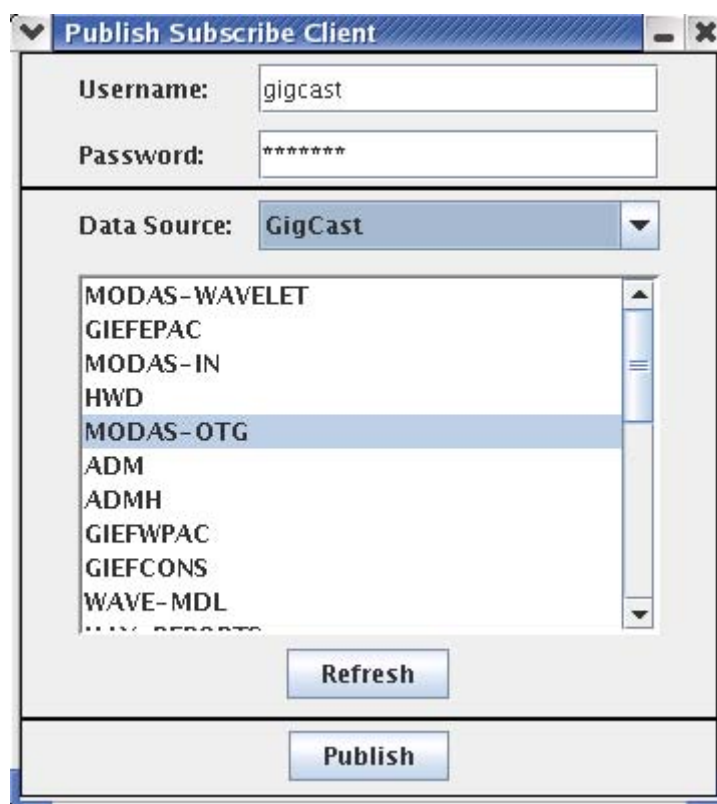
Tab B: How To Publish To A GIGCast Channel

1) Obtain the Java Publisher Client

The Java Publisher Client is available on GIG-Lite. The client comes as a tar file that includes a Java jar file, an associated DTD file and an executable script called publisher.sh. Untar the client tar file into a working directory.

2) Run Publisher

From within the working directory, execute “./publisher.sh”. The main GUI will appear.



If authentication is needed, enter your username and password in the spaces provided and hit “Refresh”. The available channels will be listed in the scroll pane.

3) Choose Channel

Highlight the channel to which you want to publish and press “Publish”. The Publisher GUI will appear.

Gigcast Publisher

Username: gigcast

Data Type: MODAS-OTG

Metadata: (* indicates required field)

=mime-type	text/otg
* institution	
* analysis-time	
* bounding-box	
* area	

File: ...

Publish **Clear**

4) Fill in Metadata

The metadata associated with the channel will be listed in the scroll pane. Metadata text fields that are grey are not editable and are considered to be fixed values by the server. Metadata labels that have an asterisk "*" are required fields.

Gigcast Publisher

Username: gigcast

Data Type: MODAS-OTG

Metadata: (* indicates required field)

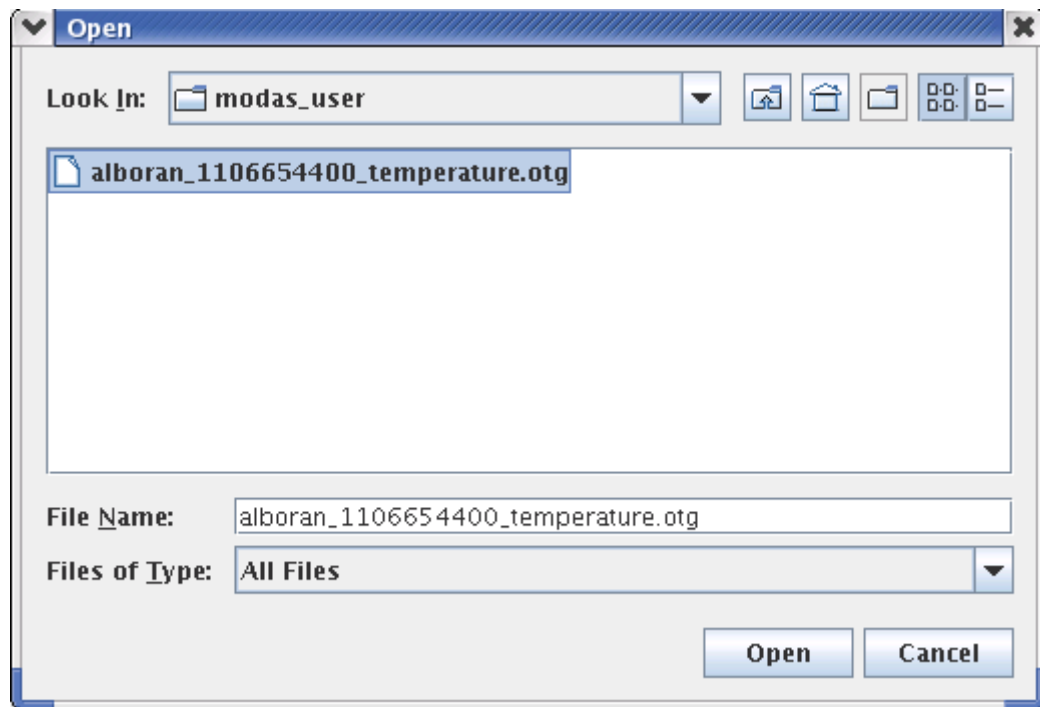
* institution	FNMOG
* analysis-time	1106654400
* bounding-box	37.0 -6.0 35.0 0.0
* area	ALBORAN
* parameter	temperature

File: ...

Publish **Clear**

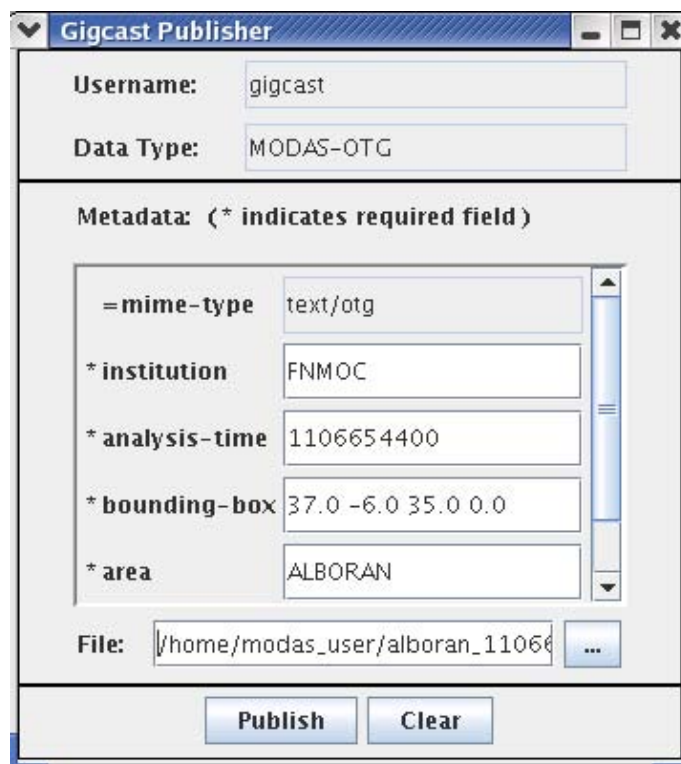
5) Choose Data File

If you know the absolute path to your data file, you can type it directly into the text field labeled “File:”. Otherwise, you can hit the “...” button and a File Chooser GUI will appear.




6) Publish

When you have filled in the metadata and selected a data file, hit the “Publish” button.



The server will send back the results.



TAB C: Mission Scenarios and Threads

You Are There Workshop: Field Intelligence

The Objective: Provide cross-coalition access to vital intelligence by a warrior on the ground.

Scenario: A Joint Special Operations Task Force (JSOTF) is operating in a hostile environment supported by a military intelligence team. The intelligence team is tasked to continuously obtain and evaluate sensor intelligence data, assess threat and opportunity, and share results with staff and operational units as appropriate. 2 X Predator and 1 X JSTARS Unmanned Aerial Vehicles are shared theater assets whose services can be requested by JSOTF. The team communicates over a 50KPS SECRET-high circuit.

Sample Value Proposition: Increase the lethality of the kill chain by breaking down the administrative barriers separating warriors from advantageous information, regardless of classification or who holds it.

You Are There Workshop: Humanitarian Disaster Relief

The Objective: Execute sense and respond logistics and command and control in support of third world disaster relief.

Scenario: After a large-scale natural disaster (earthquake coincident with heavy monsoon rains) in SE Asia, a humanitarian effort is undertaken to provide relief and stability in a devastated and remote region of a mountainous country. The disaster has eliminated roads and airfields used for accessing the backcountry. The government of the country has requested aid, and will permit US military forces into its borders to assist with initial relief efforts. Additionally, non-government relief organizations are rallying to the cause and are being permitted to enter the country.

Sample Value Proposition: Increase the speed of support to chaotic zones by employing intelligent agents against rapidly accumulating raw data to accelerate evaluation of potential courses of action.

You Are There Workshop: Border Control

Objective: Establish international border control.

Scenario: The international intelligence community reports that Al Qaeda "Chatter" is high. It is height of European tourist season and the Euro is very strong. Airports are thronged. The European Union, the United States, and most of the members of the United Nations have agreed to collaborate with respect to sharing data that might help identify terrorists at border check points.

Sample value proposition: Prevent terrorist movement by cross referencing distributed biometrics, stolen documentation, wanted persons, etc. data bases, in real time and in alignment with the myriad international agreements governing behavior.

11. Documentation of Trusted Authorization (Role Based) Policy Engine

Building Multilevel Secure Web Services-Based Components for the Global Information Grid

From *CrossTalk*, The Journal of Defense Software Engineering, May 2006

By

Dylan McNamee , Galois Connections, Inc.
CDR Scott Heller , Program Executive Office C4I and Space
Dave Huff , Fleet Numerical Meteorology and Oceanographic Center

A consensus is growing that the Department of Defense's vision of a future Global Information Grid will be built using architecture that takes advantage of Web services and uses standard Internet protocols, interchangeable components, and commercially available hardware and software wherever possible. This article describes the features and architecture of two systems: the Trusted Services Engine and the Multilevel Document Collaboration Server, including their use of a separation kernel with multiple independent levels of security, the design and assurance architecture of the cross-domain block-access controller, and the composition architecture that extends the inter-level isolation property from the block access controller outward through complex services.

The Global Information Grid (GIG) is the overall architecture intended to replace current stovepipe information systems. A consensus is growing that the Department of Defense's vision of this future GIG will use an architecture that takes advantage of Web services and uses standard Internet protocols, interchangeable components, and commercially available hardware and software wherever possible. By adopting modern standards-based protocols, the GIG will enhance current capability by enabling people and components to work together dynamically with integrated data.

Protocols such as Hypertext Transfer Protocol, eXtensible Markup Language (XML), Web-based Distributed Authoring and Versioning (WebDAV), Really Simple Syndication, and Lightweight Directory Access Protocol allow the GIG to be made of off-the-shelf components where appropriate. Where custom components are required, pervasive use of these protocols preserves the component-based architecture of the GIG, thus protecting the architecture from developing into a stovepipe system.

Many of these components and protocols are mature and well understood, but they were not designed with security as the paramount consideration. Securing the GIG is therefore a significant challenge. Particularly critical is securing its cross-domain services. For these, the GIG itself must somehow enforce separate levels of security.

Today, physical isolation enforces separation, though other technologies such as cryptography may someday be used. Such separation allows the use of commercial components as single-level components not responsible for cross-domain security concerns. However, for the GIG to realize its potential, some components must enable secure *cross-domain* data access. Clearly such components, while they must conform to commercial protocols, must be developed to *higher than* commercial standards.

This article, which describes such a component, has three main parts:

1. We describe the security and assurance attributes required of a cross-domain component of the GIG.
2. We describe the architecture and technologies we are using to achieve these attributes in the Trusted Services Engine (TSE), a network-enabled file store with integrated read-down across security domains.
3. We conclude by describing a system built on the TSE, the Multilevel Document Collaboration Server, to enable cross-domain collaboration *withindocuments* – an example of using simple cross-domain components to build more complex cross-domain systems using only standard protocols and APIs.

This article describes the features and architecture of both systems:

- The design and assurance architecture of the cross-domain *block access controller* (BAC).
- The use of a Multiple Independent Levels of Security (MILS) separation kernel.
- The composition architecture that extends the cross-domain isolation property from the MILS separation kernel to the BAC and outward through complex services.

This article is focused toward a technical audience familiar with Web services.

Assurance Requirements for Cross-Domain GIG Components

The nature and mission of the GIG makes it a prime target for trained, well-funded, and resourceful adversaries. The threats posed by such adversaries, coupled with the value of the information on the GIG, require us to show that the GIG components are robust in the face of these threats. In particular, the greater security risks associated with cross-domain components – as compared to single-level, commercial solutions – require a correspondingly higher level of trust. The process of generating and evaluating evidence of trustworthiness is known as *assurance*, the most difficult aspect of security engineering.

Two processes in the defense and intelligence communities support each other to generate assurance evidence for a GIG component: evaluation and certification. *Evaluation* is the process of validating security claims for a particular component. For example, the Common Criteria is an international standard for specifying claims of system security functionality and generating assurance that these claims are satisfied. We have determined that the cross-domain components we are building will need to meet the requirements for Common Criteria's Evaluation Assurance Level 6 or 7 [1].

Certification focuses on verifying that a component can be securely deployed at a particular site. Certification is best represented by such processes as Secret and Below Interoperability, and Top Secret and Below Interoperability. What these processes have in common is a way to tailor requirements for evaluation or certification of the following:

- Sensitivity of the data that the component handles.
- Severity of the threats it must withstand.

For example, under Director of Central Intelligence Directive 6/3, a cross-domain component that needs to demonstrate high assurance with respect to confidentiality must

satisfy Protection Level 4 or 5 assurance requirements. Evaluating or certifying a component to one of those standards requires an extensive investment in time and resources. But given the responsibilities of a cross-domain component of the GIG, high assurance is a must.

Architecture for a High-Assurance GIG Component

The TSE, a government off-the-shelf software development project funded by the Space and Naval Warfare Systems Command (SPAWAR) and National Security Agency, is a network-enabled file store with integrated read-down across security domains. The TSE provides the file store using the standard WebDAV protocol. It has a separate hardware network interface for each network security level and a separate file store for data at each level.

The TSE enforces the Bell-LaPadula policy of information flow [2], in which users on each network can read from their own level and below, but can write only to their own level. For example, when one security level dominates another (for example, TOP SECRET dominates SECRET), the TSE allows *read-down* – the ability for users at a higher level to access data from a lower level, but not vice-versa. All levels share a single name space, but views of that name space differ according to the network security level accessing the TSE.

Read-down eliminates the need for low-security data to be explicitly copied for users at high security. The single name space combined with read-down makes a wide range of applications and user workflows easier, more dynamic, and less error-prone than existing solutions.

Developing, certifying, and evaluating a high assurance cross-domain component such as the TSE at acceptable cost requires a fundamentally different architecture from that of typical, single-level components. Our approach is the following: Use as few high-assurance components as possible, each with a single purpose, to keep it small and simple, allowing it to be analyzed formally. But security is a property of a whole system, not just a component. Appropriate composition techniques can extend the security properties of the trusted computing base outward to the rest of the system.

The TSE's trusted computing base consists of the minimum number of components: one. TSE functionality is decomposed into a set of single-level components and only one cross-domain component. The underlying MILS separation kernel separates components at different security levels. Each network security level has a set of clients, an authentication service, and an integrity checker (see Figure 1). Within the TSE, each network level has its own network interface card, hard drive, and software stack implementing the TSE's networking, WebDAV, and file system services.

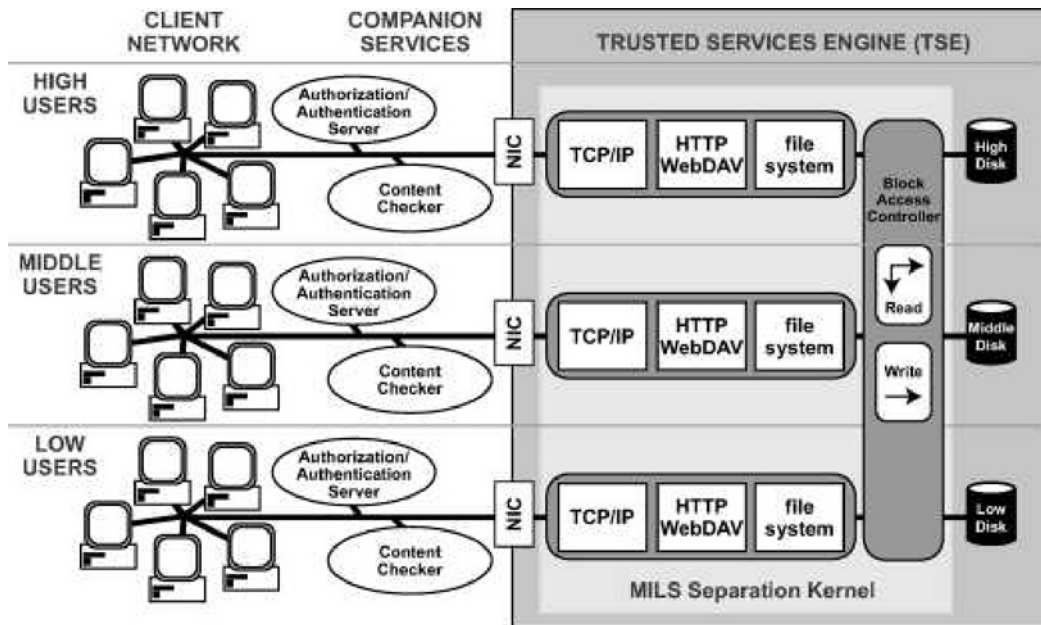


Figure 1: Trusted Services Engine (TSE) Architecture

The TSE's only cross-domain component, the BAC, mediates all access between the TSE and each level's disks.

How can these components be assembled to provide secure, cross-domain services?

1. The base must be secure before building on it. We must first establish the isolation properties of the cross-domain component.
2. We can then extend these properties to physically separate networks by mapping the software components to separate partitions in the separation kernel.
3. Finally, the separation kernel is configured to permit communication only between appropriate components.

The Cross-Domain Component

Together with the separation kernel, the BAC is responsible for isolating each level in the TSE. It is, therefore, the component that needs to be evaluated and certified to the highest levels of assurance. The BAC's functions are the following:

- Mediate all disk block access.
- Connect single-level disks and partitions.
- Write blocks to the same level.
- Read blocks from the same or lower levels.

The keys to BAC security are that it has a well-defined job and is constructed from very few lines of code. The current version of the BAC is 780 lines of C code. To ensure that the BAC implements the required attributes, we do the following:

1. Develop a formal model of the code.
2. Verify that the model corresponds to the code.
3. Develop a formal model of the policy.

4. Use model-based testing to check that the code implements the policy.
5. Formally verify that the model implements the policy.

Our formal verification ensures that the TSE security policy maps directly to the model, and the model to the implementation. To map the policy to the model, we use the Isabelle Higher Order Logic (HOL) theorem prover [3]. The theorems we prove in this logic are the following:

- None of the error states are reachable.
- The noninterference property holds.

The noninterference property states that all system actions by high security-level components are invisible to low security-level components; that is, the final state of the low-level component is the same as it would be if no actions had occurred at the high-security level.

To map the model to the implementation, a code-to-spec review team of at least two people performs a line-by-line inspection of the HOL code and the C implementation.

The example in Table 1 – a single step of the BAC – shows how closely the model matches the implementation. Our model-based testing approach uses the QuickCheck tool [4]. Based on a formal statement of the security policy, QuickCheck generates test cases that check whether or not the implementation violates that policy. The policies we have verified using this method are the following:

- **Read-across:** Reads fetch the data written at that same level.
- **Read-down:**
 - Valid reads succeed.
 - Invalid reads (that is, read-up) fail.
 - Read-downs do not affect the lower level being read (noninterference).

HOL Model	C Code
<pre> bacStep :: "config => (unit, store) m" "bacStep conf == let n = numLevels conf in processQueuedLevels (requestsPerLevel conf) n >> queueLevels conf n" </pre>	<pre> void bacStep (config conf) { nat n = conf->numLevels; processQueuedLevels (conf->requestsPerLevel, n); queueLevels(conf, n); } </pre>

Table 1: A Single Step of the Block Access Controller

Other Key Components

MILS Separation Kernel

The BAC, when hosted by the MILS separation kernel [5, 6], is an instantiation of the

reference monitor concept [7]. Unlike a traditional operating system that provides many services and abstractions, a separation kernel provides only data isolation among separate partitions and controlled communication between *partitions*. Porting an application to MILS also requires choosing a runtime or operating system to run within each partition that provides the higher-level system services the application requires, or porting one of your own choosing.

It is not enough simply to port a single-level application to a MILS separation kernel, however. The system needs to be thoughtfully decomposed and mapped to MILS partitions. Further, some key components (such as the file system) may need to be radically restructured to function in a multilevel environment.

While the TSE project aims to be portable across separation kernels, the initial target is Green Hills Software's INTEGRITY Server. This platform allows us to deploy software components from different security levels on the same hardware, thus reducing space, weight, and power requirements while retaining isolation properties equal to those provided by networks on physically separate hardware.

The WebDAV Server

The single-level components of the TSE are the WebDAV server, the file system, the network stack, and the secure sockets layer/transport-layer security (SSL/TLS). To provide the security aspects of WebDAV with high assurance, we implemented the WebDAV server using Haskell, a type-safe functional language [8]. We ported the Haskell runtime system to INTEGRITY server. The Haskell runtime system encapsulates services such as networking, threading, and memory management.

The Wait-Free File System

As Figure 1 shows, the TSE file system is a single-level component. We were surprised to find that no existing single-level file system met our requirements. The problem is caused by read-down – a user on a high-level network can read files from a lower level while a user on the low-level network changes those files. Ordinarily, locks could be used to solve this problem, but cross-domain locks violate non-interference and are unacceptable in this case. How can the TSE present consistent data without introducing a proscribed communication channel, overt or covert?

Designers of algorithms for shared-memory multiprocessors face a similar problem that they solve using a method called *wait-free synchronization* [9]. Wait-free synchronization guarantees that interactions with concurrent objects take a finite number of steps instead of using *critical sections*, which block competing processes for an indeterminate time. The *wait-free file system* adapts this idea for its own synchronization method. This preserves the isolation property by the following:

- Writers are oblivious to readers.
- Readers can proceed independently of writers.

Outside Services

To minimize the trusted base and avoid duplication of function, the TSE will use, or uses outside services wherever possible. Key services are authentication and integrity-checking; so far we have evaluated Navy enterprise single sign-on for authentication and

one-way file transfer for integrity-checking, but final decisions will be driven by the demands of specific installations at customer sites.

Though it is conservative and efficient to draw on outside services, it also means that we must build a chain of trust from our base to the outside service. We use several methods to help us do so:

- Outside services are all single-level, which minimizes their trustworthiness requirements.
- We choose services specified and trusted by our customers that have been vetted in similar deployment scenarios.
- The TSE and companion services use the standard cryptographic protocols SSL/TLS and digital certificates to manage communication between them.

The sum of the TSE and a specific set of external services is submitted for the certification prerequisite to multilevel deployment.

Building Complex Multilevel Services on the TSE

The TSE can be used as a building block for more complex cross-domain services, as demonstrated by another current Galois project, the Multilevel Document Collaboration Server (DocServer). Its architecture reuses the decomposition structure of the TSE to provide multilevel secure document-based collaboration.

The DocServer allows a user at a high network level to make private modifications to an XML-based document stored at a lower level. The DocServer supports ongoing modifications at multiple network levels; modifications from the high network are visible only to users on the high network, while modifications from the low network are visible to users at that level and above.

The DocServer also supports publishing regraded documents from high network levels to low, using XML filtering and integration with an outside regrading system such as Radiant Mercury or ISSE Guard. These systems enable transfer of documents from high security to low security by enabling a human reviewer to reliably review all of a document's contents (including possibly hidden content), and, upon successful review, write it to the low network.

In the case of the DocServer, a high-level user marks up the document according to a new set of security levels, and submits it for regrading. The DocServer filters the document and sends the filtered version to the regrading system. After human review, the filtered version of the document is written to the DocServer's low-level file system.

Figure 2 shows the *publish, edit, merge* workflow of the DocServer. At left, a user on the Secret network publishes the document to the Unclassified network. The DocServer filters the Secret content and submits the resulting unclassified document to the regrader. After regrading, users on both network levels make modifications to the document. Modifications made at Secret are not visible below, but Unclassified modifications are visible to users at Secret using the DocServer's *merge* each time the document is read.

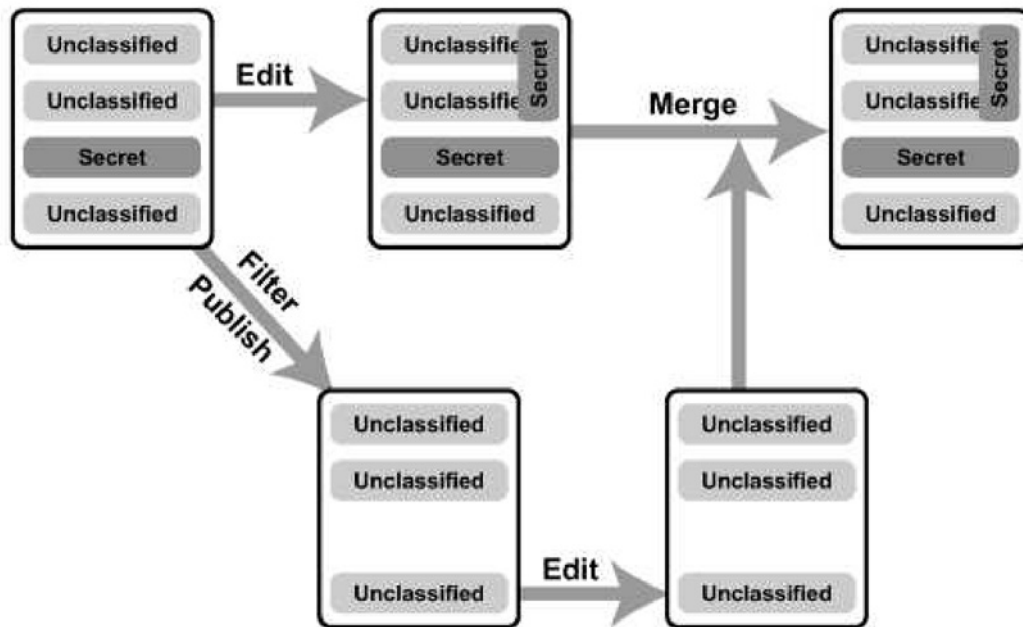


Figure 2: DocServer Merge Operations

The DocServer is a Phase 1 Small Business Innovative Research project funded by SPAWAR.

Conclusion

The DocServer uses the TSE for file storage and its sole cross-domain component. Reusing the only high-assurance component gains us a great deal – the DocServer should be certifiable to the same level as the TSE with little additional work.

The DocServer's use of the TSE to achieve high assurance, cross-domain function mirrors the TSE's internal use of the BAC. By building the DocServer from this core component, we once again take advantage of the BAC, effectively extending its security policy through to increasingly complex systems.

The TSE's component architecture demonstrates a powerful technique for extending the security properties of a formally analyzed core component to a wide scope. In a similar manner, the DocServer uses MILS to extend the security properties of the TSE outward to provide complex multilevel functionality.

TSE Status

Development of Vers. 1.0 of the TSE will be complete in summer 2006, and will be followed by certification at a customer site. We expect to begin Common Criteria evaluation at evaluation Level 6+ the following year. Phase 1 of the DocServer is near completion. We hope to begin Phase II in spring 2006, and commercial transition sometime in 2007.

Acknowledgements

The authors would like to acknowledge contributions from the following people: David Burke with the evaluation and certification sections; John Matthews and Paul Graunke with the verification and validation sections; and Lauren Ruth Wiener with the clarity of thought and exposition.

References

1. Common Criteria <www.common-criteria-portal.org>.
2. D. E. Bell and L. J. LaPadula. *Secure Computer Systems: Mathematical Foundations and Model*. The Mitre Corporation, 1976. <<http://csrc.nist.gov/publications/history/bell76.pdf>>
3. Isabelle. “A Proof Assistant for Higher-Order Logic.” University of Cambridge Computer Laboratory <[www.cl.cam.ac.uk/Research/HVG/ Isabelle](http://www.cl.cam.ac.uk/Research/HVG/Isabelle/)>.
4. QuickCheck Automatic Specification-Based Testing <www.cs.chalmers.se/~rjmh/QuickCheck>.
5. National Information Assurance Partnership. *U.S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness*. Vers. 0.621. Ft. Meade, MD: NIAP, July 2004 <http://niap.nist.gov/pp/draft_pps/pp_draft_skpp_hr_v0.621.html>.
6. Vanfleet, Mark W., et al. “MILS: Architecture for High-Assurance Embedded Computing.” Aug. 2005 <www.stsc.hill.af.mil/crosstalk/2005/08/0508vanfleet_etal.html>.
7. Anderson, James P. “Computer Security Technology Planning Study.” Fort Washington, PA: James Anderson & Co, Oct. 1972 <<http://csrc.nist.gov/publications/history/ande72.pdf>>.
8. Haskell. Haskell: A General-Purpose Purely Functional Language <www.haskell.org>.
9. Herlihy, Maurice. “Wait-Free Synchronization.” *ACM Transactions on Programming Languages and Systems (TOPLAS)* 1: 124-149. New York: ACM Press, Jan. 1991 <<http://portal.acm.org/citation.cfm?id=102808>>.

About the Authors



Dylan McNamee, Ph.D., is the technical lead for cross domain projects at Galois Connections. He received his doctorate in computer science from the University of Washington.

Galois Connections
12725 SW Millikan WY
STE 290
Beaverton, OR 97005
Phone: (503) 626-6616 x137
E-mail: dylan@galois.com



CDR Scott Heller is currently the Cross Domain Solutions lead at PMW 160 within the Program Executive Office Command, Control, Communications, Computers, and Intelligence, and Space in San Diego, Calif. He has a master's degree in computer science with an emphasis in Multi-level Security from the Naval Post-Graduate School in Monterey, Calif.

Program Executive Office
C41 and Space
626 Orange AVE #303
Coronado, CA 92118
Phone: (619) 929-1451
E-mail: scott.heller@navy.mil



Dave Huff serves as the director, Exploratory Projects Division at the Fleet Numerical Meteorology and Oceanography Center. His team is focused on information assurance and Web-based techniques for establishing identity, authorization, and cross-domain information exchange.

Fleet Numerical Meteorology and Oceanographic Center
7 Grace Hopper AVE
Monterey, CA 93943
Phone: (831) 656-4569
E-mail: dave.huff@metnet.navy.mil

12. Netcentric Certification Office Statement of Work

Statement of Work - Joint Interoperability Test Command (JITC) Netcentric Certification Office (NCO)

Principle Investigator: C. R. Gunderson, Department of Information Science, and the Cebrowski Institute, Naval Postgraduate School

Period of Performance: 1 April 06 – 31 Dec 06

Background and Approach

Today's test and evaluation process for C4ISR acquisition programs is *system centric*. That is, we formally evaluate whether a particular combination of hardware and software meets "system performance specifications" in formal developmental and operational test environments before fielding the stand alone "system". This is a long, expensive, one-size-fits-all, serial process that is out of touch with the rapid information processing technology refresh cycle.

Conversely, Global Information Grid (GIG) Netcentric Enterprise Service (NCES) vision calls for deploying a system of services, wherein distributed composable software components will be fielded rapidly and piecemeal, in deliberate expectation that the services will be "discovered" and applied netcentrically in unpredictable and uncontrolled ways. To realize this vision, the Director of DISA has mandated an Adapt, Buy, Create (ABC) policy that calls for first Adapting existing tools, or after determining that existing tools are inadequate, Buying generic commercial tools and adapting them to solve specific problems, or as a last resort, Creating specialized solutions. Emphasis is squarely on innovative re-use of existing components.

Accordingly, rather than a traditional closed "system," DISA intends that NCES and the associated next generation Joint Command and Control capability, JC2, will be a highly distributed *System of Services* composed primarily of best of breed existing components, select COTS components, and a relatively small number of created specialized components. Barriers to this approach include current lack of incentive to encourage and process to allow program managers to perform adaptation across stove-pipe development domains, and lack of a top down systems engineering perspective on DoD enterprise-level network capability requirements, resources, and gaps upon which to base such incentives and process.

Clearly, achieving the objectives of GIG NCES in general, and JC2 in particular, will require a fundamentally new way to perform development and T&E, i.e. an *information-centric* approach. Discoveries made in the W2COG Research Initiative suggest an approach. Government and industry *NCES and JC2 developers and customers should partner as peers employing best of breed e-business methodology to*

concurrently, develop, evaluate and field netcentric capability via consumable off the shelf model. Accordingly, we propose an agile and pragmatic risk/reward-based process that literally teams operational users with JC2 and NCES program managers, developers, and testers, in rapid spirals formed around bundling off the shelf net-enabling components to impact real-world operational imperatives. These spirals will address policy, doctrine, and CONOP as well as technology. We will employ both top down enterprise-wide capability and process analysis and bottom up validation and verification of netcentric productivity enhancements targeted and achieved. We will quantitatively balance the need to protect information with the need to share it, and introduce the notion of *value per bit exchanged* into the test process. We will employ the Federated Development and Certification Environment (FDCE) concept by linking the DoD's distributed test facilities into a virtual software test range we call the Netcentric Certification Office (NCO). Test and evaluation strategies will flex according to deliberately considered technical and operational risk versus operational urgency. These strategies will include notions of rigorous "type certifying" classes of software components as "well behaved network citizens" and subsequent "accrediting" of compliance via relatively light weight auditing techniques. Consequently, we will push most of the responsibility for developmental testing to the software developers themselves, and use operational testing to verify capability enhancement achieved by bundling components. Successful net-enabling reference implementations will be fully documented and made immediately and broadly available via consumable off the shelf model. We propose an initial research pilot to demonstrate the viability of this approach followed by rapid operationalization.

Statement of Work for JTIC

The Naval Postgraduate School Cebrowski Institute will leverage lessons learned in its W2COG research initiative and apply expertise resident among Naval Postgraduate School Faculty and Students to assist JTIC create the Netcentric Certification Office and thereby accelerate incremental fielding of NCES and JC2. In particular, the Cebrowski Institute will assist JTIC to design pilots and follow-on operational instantiations of to the following tasks:

1. Coordinate activity of existing distributed developmental, test, and evaluation assets rather than create new infrastructure or bureaucracy.
2. Perform DoD enterprise level netcentric capability and process analysis.
3. Develop a standing contractual vehicle between JTIC and the W2COG Institute to allow the DoD to partner directly in joint piloting "ventures"

with a broad spectrum of industry, government, and academic information processing expert operators and developers.

4. Form “ecosystems” of operators, vendors, labs, and sponsors around compelling operational issues, mature technology, and ~90 day spirals to demonstrate, prototype, and productize bundled information processing capability and evaluate potential adjustments to associated doctrine, policy, and CONOPS.
5. Assist operational units and other customers develop netcentric productivity metrics, i.e. measurable increase in the value of bits netcentrically available to operators as a result of a new or improved tool. Netcentric productivity targets will form the basis of requirement validation.
6. Coordinate authorizing, scheduling, and funding for commercial use of DoD network test facilities.
7. Test according to the following criteria:
 - a. Compliance with/or utility of, existing or proposed SOA/GIG technical reference, policy, and/or doctrine.
 - b. Alignment with COTS standards, trends, and investments
 - c. Transition risk assessment including schedule
 - d. Utility in mission context (e.g. power requirements, user friendliness, bandwidth reality, durability, size/weight)
 - e. Performance verification with respect to netcentric productivity metrics.
8. Certify and/or accredit that off-the-shelf network-enabling software satisfies government requirements including documentation of reference implementation details: hardware, software, and interface specifications; training and doctrine requirements.
9. Place certified and/or accredited software on approved consumable procurement schedules.
10. Establish and maintain a “GIG-lite” distributed on-line repository of off-the-shelf network enabling software (both certified and pending certification) and mission thread based modeling and simulation demonstration suite.
11. Maintain open source library of contributed code generated in NCO activities.


13. Selected Media Clippings

Consortium Jump Starts Network-Centric Interoperability

September 1, 2005

by Capt. Chris Gunderson, USN (Ret.)

Semantic interoperability is an age-old problem that has confronted mankind since the Tower of Babel. Finding a solution to the confusion of tongues is the next plateau in the U.S. Defense Department's Global Information Grid strategy. In pursuit of that modern day Rosetta stone, an organization comprising representatives from the military, government, industry and academia has created an open forum for technical experts to brainstorm and rapidly field interoperable information-processing solutions across Defense Department and coalition partner networks.



The [World Wide Consortium for the Grid](#) (W2COG) is the brainchild of Peter Denning, professor and computer science department chairman, Naval Postgraduate School, Monterey, California. It is a working meritocracy based on the principle that networks of motivated experts, independent of central authority, are the best and only way to solve complex problems, Denning says. "The focus is operational, on real-world mission-thread-analysis-driven communications-architecture solutions to top-down Department of Defense Global Information Grid policy," he explains.

"There are two major challenges to achieving robust interoperability across distributed networks—one is technical and the other is organizational," Denning maintains. "The technical challenge is how to put together the best information environment to speed and ease information sharing among systems that were not originally designed to talk to one another. The organizational challenge is determining how to get bureaucracies to adopt new practices that cut across organizational boundaries and harness collective expertise. The W2COG simultaneously addresses both of these challenges."

To facilitate evaluation and comparison of operational network-centric solutions, the W2COG's Web site will host a portal to a credible virtual simulation of the enterprise architecture envisioned for the GIG. The simulation is called GIG-Lite. Working engineers and technical experts on the GIG domain and network-centric operations are encouraged to join the consortium via the site. All ideas and solutions members propose will be quickly evaluated and selected based on merit for informal and formal field-testing by technical peers.

To launch the new consortium, the [Naval Postgraduate School](#) jointly sponsored an inaugural W2COG Working Symposium in May. Other symposium sponsors included George Mason University, the Network Centric Operations Industry Consortium, the Association for Enterprise Integration and the Association for Computing Machinery. More than 130 network-centric experts from the active duty and retired military, civilian agencies, coalition allies, industry, academia and nongovernmental organizations participated.

Conference attendees identified six engineering pilot projects to advance the consortium's mission, including one that leverages portable hastily formed networks developed at the Navy's corporate university. As a result of the success of the first symposium, a second is planned for February 2006.

The W2COG's founding team was inspired by the World Wide Web Consortium (W3C), a key catalyst for Web technological innovation that currently has more than 400 member organizations. Don Brutzman, associate professor of information sciences, Naval Postgraduate School, and founding W2COG team member, is the W3C's liaison to another key group, the Web3D Consortium (Web3D). Web3D is developing the Extensible 3D (X3D) graphics specification for three-dimensional graphics on the Web, a project in which Naval Postgraduate School faculty and students are playing major roles.

"We're now at an exciting threshold to the GIG," Brutzman relates. "The Rosetta stone for getting data to semantically interoperate and structuring information for use on the Web is XML [extensible markup language], a meta-language for writing other computer languages; and interoperability using XML, in turn, is the heart of the Defense Department's Global Information Grid strategy. So we now have all the pieces to do network-centric operations—they just need to be put into place. That's the reason for W2COG.

"We're taking full advantage of past successes in deciding how to set up W2COG," Brutzman notes. "The Internet gave us data communications; the Web made any raw data available; and the third plateau, which we're now working on, is semantic interoperability. Once that's in place, truly new and exciting Web capabilities can emerge. W2COG's early members will be those who want their organizations to communicate using XML, the GIG and Web services. That's a big step, but it's the future of all defense, government and intergovernmental communications—working together like never before."

The Office of Force Transformation; the Assistant Secretary of Defense, Networking and Information Integration; the Undersecretary of Defense, Advanced Systems and Concepts; the Defense Advanced Research Projects Agency; and the Space and Naval Warfare Systems Command were the Defense Department sponsors for the W2COG's planning phase. The Naval Postgraduate School will take the lead in coordinating research that addresses W2COG-identified topics and will provide access to its large pool of master's degree students who conduct thesis research.

Denning says that a long-term W2COG goal is to develop a prioritized agenda of technical questions that need to be researched. "These will be matched with Naval Postgraduate School student thesis topics through a formal, coordinated process, as well as coordinated with the wider university community," he explains.

The professor has asked Maj. Angela Burth, USAF, communications officer, to help create and lead the W2COG's market/broker solutions working group. "The main goal of the GIG is interoperable information sharing among a large number of often ad hoc partners, and markets are the planned transaction space for such a practiced adhocracy," Maj. Burth says. "One means to achieve this goal is using market mechanisms to stimulate and facilitate information sharing. In fact, W2COG itself is a network-centric solution, e-transaction space, which is why it's a true force for change."

The W2COG is forming strategic partnerships with two related industry consortia: the Association for Enterprise Integration (AFEI) and the Network Centric Operations Industry Consortium (NCOIC). The AFEI will concentrate on policy, while the NCOIC will focus on engineering standards and the W2COG will work on building and testing prototypes.

Capt. Chris Gunderson, USN (Ret.), is an associate research professor of information sciences at the Naval Postgraduate School, Monterey, California, and the executive director of the World Wide Consortium for the Grid. He can be contacted at chris.gunderson@w2cog.org.

Netcentric Warfare Looks Like Commercial E-Business. Just Ask Al Qaeda.

By Chris Gunderson.

Successful enterprises understand their business model and use technology to support it. The model comes first.

Let's say you are the boss of a multinational outfit that owns a giant computer network made up of hundreds of independent systems built, operated and maintained by scores of subordinate organizations. You've made new business decisions around the idea of globally distributed operations. Wisely; after all, that strategy works for Wal-Mart, FedEx, and many others.

So, you task your heretofore autonomous subordinates with pooling their resources and building a global computer network that allows interoperable connectivity with legacy and ever-increasing new data sources and network services. Obviously, you want to encourage these subordinates to continuously improve the interoperability and accessibility of the connected data sources and services. Name a company with success at these kinds of tasks. Here's a hint: Rhymes with "bugle." Name a giant which has, so far, failed. No hints.

Meanwhile, let's say your business rival is beating the pants off you when it comes to distributed networked operations. Shockingly, your rival is a tiny little outfit that doesn't even own its own network infrastructure. It just buys services.

This isn't hypothetical. The big outfit is the Defense Department, with its visionary network known as the Global Information Grid (GIG).

The rival is Al Qaeda. Al Qaeda uses the Internet, common wireless technologies, cell phones and commercial security packages to coordinate its extraordinarily effective operations. By first hand accounts I've heard, they're better at it than we are. Any credible approach to fielding the GIG must turn this asymmetric disadvantage into an overwhelmingly symmetric advantage.

Immediately.

We've got more money, people and eye-watering technology. We should rule the information-processing marketplace, but we don't.

So, there are lessons to be learned from e-business service providers like America Online, Yahoo, and Google, and successful e-business service consumers like eBay, Amazon.com and Travelocity. This is not a surprising suggestion – especially from me. After all, I am a researcher sponsored by Office of the Secretary of Defense to study the Internet community. Neither is this a unique idea. Many in the DOD are already getting good-enough commercial technology fielded for military applications.

What I am suggesting is not that DOD merely learn lessons from the best practitioners of e-business. Rather, I am suggesting that DOD literally join the e-business ecosystem. I propose that DOD harvest the benefit of technology as it matures using the same approach and realize the same economies and speeds enjoyed by successful Internet technology consumers.

Conventional wisdom says DOD requirements are just too unique for generic solutions, but conventional wisdom is an anathema to transformation. And conventional wisdom isn't helping us beat the insurgency.

What I am suggesting requires a transformation. DOD must change from being a builder and operator of networks and become an enabler and manager of network services that will give our soldiers information advantage over our adversaries in the way Wal-Mart overwhelms its competition.

The transition has five steps.

First, develop a business model and two business targets. The business targets are to decrease network costs by leveraging economy of scale as a peer of Google, AOL, Yahoo and others. And to reinvest the savings in "innovation" just as Google et al do. DOD's objective, of course, is to gain "information edge" over an adversary.

Second, structure contracts as e-service consumers, not as program managers. Govern and create incentives for improved metrics in quality of service metrics, security, productivity and cost rather than system specifications.

Third, select GIG service providers based on success in e-business rather than defense systems.

Fourth, create test and evaluation methodologies that get network services into the field faster, then continuously improve them. Listen to user suggestions for improvements.

Fifth, align DOD engineering processes and human resources with the new mission model. Make auditing and budgeting tools match these transformational objectives. Converting legacy process and people to the new model will require a sustained sense of urgency and a relentless incremental approach.

If it is not a big change -- fast -- it is not transformation. If it doesn't help us win, it's not worth it.

Chris Gunderson is executive director of the Worldwide Consortium for the Grid (W2COG) Research Initiative, sponsored by the Office of the Secretary of Defense. He is also associate research professor of information science for the Naval Postgraduate School. He retired as captain following 27 years as an oceanographer for the Navy. Contact him at Chris.Gunderson@w2cog.org.

DoD Turns to Industry for the Internet it Wants

[GCN Home](#) > [04/03/06 issue](#)

By William Jackson, GCN Staff

The Defense Department recognizes Version 6 of the Internet Protocols as central to its concept of network-centric warfare. But enabling a worldwide network to pass IPv6 packets is not enough to realize its goal. It requires applications and tools.

This is the job of the World Wide Consortium for the Grid.

“The DOD was taken by surprise by IPv4,” said W2COG executive director Chris Gunderson.

Commercial products developed to open technical standards helped the Internet develop in ways never envisioned by the Defense Advanced Research Projects Agency. DOD wanted to leverage that power for its Global Information Grid.

But “leadership realized that DOD wasn’t wired that way,” Gunderson said. “We continue to develop the same capabilities with different vendors again and again and again. That’s why [DOD] felt it needed to invest in an organization like this.”

W2COG was formed as a commercial incubator in late 2004 with \$1.6 million in DOD seed money and was incorporated in June 2005. It helps the department work with academia and the private sector to develop commercial products for its own net-centric operations.

Is it too late for DOD to influence the commercial development of IPv6 for its own ends?

“Nobody is very far along in kicking the tires on it and figuring out what it can do,” Gunderson said. “There is still plenty of time for DOD to be on the leading edge.”

W2COG still is ramping up. “We have only now started to collect dues and get our first products on the table.”

The first product making its way through the pipeline is an ultrawideband, wall-penetrating radar device that could create a security “bubble” for troops in the field. When an area, building or room has been cleared of combatants, a device could be installed that would alert troops if anyone re-entered the area.

The Marine Corps Systems Command is funding testing and development of the device, which now is in the demo stage. “By May it will be a product in the field,” Gunderson said.

W2COG in January announced a partnership with the IPv6 Forum to promote and provide resources for development and deployment of next-generation networking technologies.

“IPv6 gives us the opportunity to improve current communications capabilities, especially at the edge of the network,” where it is needed most in DOD net-centric operations, Gunderson said.

Incubators such as W2COG are helpful in developing new technologies, because established companies often are reluctant to take these risks.

“The small companies hungering at the edge are going to come up with these ideas,” Gunderson said. “What we’re providing is a place where you can fail fast and cheap,” so successful ideas can quickly be sifted out and developed.—William Jackson

14. Valuable Information at the Right Time (VIRT) and Value off the Shelf (VOTS) A Formula for Netcentric Engineering

**A Market-based Approach to Building the Global Information Grid:
Valued Information at the Right Time and
Value off the Shelf ... VIRT & VOTS!**

Executive Summary

The US DoD intends its enterprise architecture, the Global Information Grid, to deliver high-value services aimed at two key objectives: enhance mission effectiveness through more efficient information processing, and enhance programmatic efficiency through more efficient acquisition process. Both depend on the concept of netcentric operations (NCO), i.e. *effective* interaction of distributed, composable, modular components. “Effective” means measurably improves netcentric productivity, where “productivity” is defined as effect/effort. Hence we must focus on how to improve the value of the information that NCO processes manipulate.

Therefore, we should pursue a market-based approach to defining information value. Call this operational concept "Valuable Information at the Right Time" (VIRT) and the acquisition concept "Value Off The Shelf" (VOTS). Communities of interest (COI) will define the value attributes of VIRT and VOTS, i.e. what DoD has coined “net-ready” key performance parameters (NR-KPP). COIs will require new doctrine and methodologies to accommodate this new concept. W2COG proposes that we adapt the “Internet model” of enterprise engineering -- which we call an *NCO Incubator* -- to rapidly field information processing *reference implementations*. A reference implementation is a successfully demonstrated instance of incrementally improved information processing capability, one increment of NCO productivity. In addition to reference implementations, the *NCO Incubator* will deliver documented lessons learned with respect to both reference implementation technology and supporting COI process and doctrine.

The *NCO Incubator* differs from existing government rapid prototyping efforts in an important way: vendor innovation and internal R&D are leveraged up front and the government will work to place successfully demonstrated products immediately on the GSA or other convenient consumable procurement schedule. Thus the implementation of the NCO incubator requires a component inside government and a component outside, working together.

The W2COG Institute, with its network of government and industry experts, is the non-government component of the NCO Incubator. We propose that the government establish a pilot project to demonstrate in-government component, which we refer to as the Unifying Netcentric Program Office (UN-Program Office!)

A Market-Based Approach to Defining Information Value

DoD's Global Information Grid (GIG) is an approach to service-oriented enterprise architecture that aims to improve mission effectiveness and programmatic cost effectiveness. The key to achieving both goals is net-centricity, i.e., effective distributed collaboration among composable, interoperable, components. Adopting netcentric operations and engineering represents profound change in DoD practice, i.e. "transformation." Fielding GIG will require the following two doctrinal imperatives.

1. **Valued Information at the Right Time (VIRT):** an approach to the increasingly vast market place of data that encourages careful choices regarding how valuable time is spent "buying" and "consuming" information. We get competitive advantage from finding opportunity and acting quickly. We lose competitive advantage if we waste time processing insignificant data. Time spent acting on critical information makes money and wins battles. A decision made too late is not a useful decision.
2. **Value Off The Shelf (VOTS):** an approach to delivering "good enough" capability fast enough to take advantage of the rapid advances taking place in the information technology market. This approach emphasizes bundling interoperable off-the-shelf components in an architecture optimized for VIRT. We get competitive advantage from finding opportunity and acting quickly. We lose our advantage if we waste time on processes that are inherently too slow to exploit current opportunities. Time spent fielding capability makes money and wins battles. A system fielded after the battle is not a valuable capability.

Notice that these doctrinal imperatives depend on an entrepreneurial market behavioral model. Neither existing DoD operational nor acquisition doctrine confront entrepreneurial market models. Both dictate disciplined hierarchies with associated rigid behavior models. Netcentric transformation requires change to, or at least addition of, this value-oriented approach to focusing basic behavior across all levels of the enterprise.

VIRT Architecture improves effectiveness by increasing the fraction of time spent processing valuable information. Important characteristics of VIRT-based organizations include the following:

1. Operators request "bids" from various potential providers to satisfy information requirements for specific tasks and desired outcomes.
2. A system of systems automatically, but sparingly, delivers context-sensitive information culled by intelligent agents that detect significant changes in the state of important parameters.
3. The system rewards information providers whose products tangibly enhance productivity.
4. The system pressures operators to make effective selections of information and information sources.
5. The system applies pressure to operators and information providers to reinforce and improve productive choices.

6. The system decomposes mission performance threads into information *value delivery chains*, comprising essential tasks, desired outcomes, and information exchange elements (IEE).
7. Mission thread analysis determines objective value characteristics of IEEs.
8. Mission outcomes determine productivity where “productivity” is defined as effect/effort.
9. Organizations design system specifications and associated performance metrics based on this definition of productivity
10. A process continuously implements and improves all of the above

A VOTS delivery process continually increases the productivity of employed systems by drastically reducing the time, cost, and bureaucratic friction for implementing improvements. This process has the following characteristics:

1. VOTS process delivers VIRT products and components that work immediately, “out of the box,” in standard computational environments.
2. Customers shop for VIRT products and components from various potential providers to satisfy specific information requirements for specific tasks and desired outcomes. (e.g., finding these preconfigured on the GSA schedule)
3. The system tangibly rewards providers of VIRT products and components that tangibly enhance productivity.
4. The system pressures program managers to make effective selections of VIRT-enabling components.
5. Product or component “value” depends only on its ability to enhance operators’ productivity = effect/effort.
6. Product and component specifications and associated performance metrics depend on ability to deliver IEEs in ways that demonstrably enhance value.
7. A process continuously implements and improves all of the above

Examples of VIRT architectures and VOTS delivery processes occur in the commercial marketplace. Target Stores uses one such implementation (Hewlett Packard’s version) called Zero Latency Enterprise (ZLE) to operate all its retail activities. (Picture the magnitude of Target’s retail activities!) ZLE links multiple heterogeneous legacy information systems, tracks tens of millions of daily transactions and hundreds of millions of inventory items in real time, links various independent organizations, manages terabytes of on-line and archived data interactively, assesses and integrates human factors, scales up or down, and is highly reliable 24/7. Target uses productivity metrics to identify and track key IEEs. Furthermore, they employ business rules and intelligent software agents to analyze the data continuously and to trigger human actions intended to enhance productivity. For example, fraudulent transactions drain productivity. ZLE implementers developed an agent that selects the few most relevant IEEs required to predict whether a particular transaction is fraudulent from the best sources of that data. As a cashier blithely executes what is likely to be a fraudulent transaction, the agent automatically sends an alarm to enable real-time intervention by security officers. The full-scale pilot project that launched ZLE as a product took ~90 days to execute. Upon completion of the pilot, ZLE was offered as a consumable product. Impressed by the demonstration, various retailing enterprises bought ZLE “off the shelf” and quickly scaled, tailored, and fielded it.

Communities of Interest and the Internet Model

Although VIRT architecture and VOTS sustaining and improving process are not yet prevalent in DoD, emerging GIG policy includes both ideas. The DoD Architectural Framework (DODAF) emphasizes the need to establish architectural measures of effectiveness (MOEs) based on VIRT-like principles. Further, DoD's "NET-READY" Key Performance Parameter (NR-KPP) requires that all major information systems demonstrate netcentric value added. Likewise, GIG policy intends for DoD to collaborate with industry to realize the potential of service-oriented architecture to employ re-usable, plug and play (i.e. VOTS) network service components. A goal here is to drive down the cost, and increase the speed to market of, net-enabled capability. OSD has introduced the notion of "Communities of Interest" (COI) to flesh out the detail associated with these imperatives, but we must define the specifics of such COI activity.

To field netcentric capability, COIs must understand netcentric requirements, i.e., define the specific characteristics of information "value" for a given application. The value of information depends on its ability to improve productivity. Often we struggle to define "productivity" in the context of information processing activities. One proxy for productivity is "speed to better decision." After all, "speed to better decision" correlates directly to "time not wasted evaluating irrelevant information." By working with operational units to analyze specific, critical, recurring, mission threads, COIs can quantify both "speed" and "better" in terms of the IEEs associated with particular families of recurring decisions, and thereby define the VIRT requirement.

Consider this scenario and associated mission thread. A commander wants to find and kill an elusive mobile target. The notional mission thread components are as follows: (1) receive "tipper"; (2) confirm validity; (3) evaluate risk of collateral damage; (4) evaluate and select weapon options; (5) kill target; (6) assess results. The commander might value IEEs such as unattended ground sensor reports, human intelligence reports (HUMINT), commercial satellite imagery, very high resolution classified satellite imagery, blue force location reports, etc. If the window of kill opportunity is short, timeliness becomes a value multiplier. If the target is tiny, low resolution imagery is useless. HUMINT is potentially very valuable, but only if judged reliable. Data is valueless if it is too highly classified to share with those who need it. Data requiring lengthy analysis processes will be too late to be useful. Etc. COIs can use these factors to define an information value hierarchy, and associated business rules, analogous to the way Target employs ZLE.

Netcentric Operations Pilot Process

W2COG proposes superimposing convenient, cost-effective, opportunities for COIs to rapidly field pilot reference implementations of VIRT NR-KPP specifications and MOEs while at the same time "stocking the shelves" with VOTS components. Certainly, DoD already uses the concept of pilot projects to inject innovative technology or concepts aimed at enhancing critical capability. Some of these DoD pilots are large-scale with respect to time, money, and impact. Others have more modest goals, smaller budgets, and shorter time scales and may even include COTS components. However, in every case the intended sustaining mechanism for successful demonstrations is "acquisition" via the

Program of Record (POR), i.e. the procurement system designed to field military-specific systems. A common result is that local units benefit temporarily from the “left behind” capability, but the POR fails to adopt it. Even in the cases where the new capability is embraced by the POR, it generally takes the acquisition process years to field it broadly. A capability late is thus years of productivity lost.

W2COG pilots will depart from other DoD pilot methodology by using the “Internet model.” Loosely organized participants, motivated to demonstrate interoperability, bundle their offerings together through mutually agreed open standards. Pilot projects rapidly demonstrate incrementally improved capability at low, shared, cost. The demonstrated tool(s) is called a “reference implementation” of the targeted capability. Customers help develop the tool, adopt it when demonstrated to be useful, and adjust their training and doctrine accordingly. The productized tool is put on the market, and the COI carefully studies, distills, and shares lessons learned from the reference implementation demonstration and ensuing customer feedback. An important discriminator of W2COG pilots will be emphasis on the procurement mechanism appropriate for sustaining short life cycle and broad applicability, i.e. the “consumable” model, rather than the DoD POR “acquisition” model. Specifically, W2COG will broker across traditional domain boundaries to combine various netcentric government, vendor, and academic activities. The W2COG teams will execute rapid information processing pilot projects with the following deliverables:

1. Immediate posting of successfully demonstrated VOTS components on the GSA or other convenient procurement schedule.
2. NET-READY reference implementation specifications and metrics for continuing use as “government furnished equipment” (GFE), as follows:
 - Effective COI activity
 - Hardware/software
 - COI-defined VIRT requirements
 - NCO doctrine
3. Process to empower and incentivize industry to innovate toward GIG objectives
4. Leverage of vendors’ internal R&D resources
5. Fielded capability for participating customer(s) within twelve months.
6. Virtual GIG (VGIG), a distributed web service-oriented run-time environment that simulates the GIG objective vision. VGIG will serve as an on line depository of VIRT products and components, i.e. a VOTS capability inventory. It will include mission thread simulation data bases and evaluation tools that will allow customers to evaluate offerings in proper context: a fast, low-friction delivery mechanism for VOTS products. It will also serve as an open source development environment for an interoperable service-oriented “substrate” of the GIG. .

A Netcentric Operations “Incubator”

The W2COG project has designed an “NCO Incubator” to facilitate the pilot project activity described above. This NCO Incubator will have both an industry and government component.

The industry component is a not-for-profit brokering service among agile and technically expert government and industry providers and consumers of network-enabling artifacts. Because its activities occur outside the POR, this industrial component is unconstrained by DoD acquisition regulations. It is free to choose participants and technologies on the basis of value added alone. Likewise, vendor participants in the not-for-profit industrial venture are not constrained regarding future dealings with the government. For example, the government frequently contributes resources to help develop open industrial standards in projects facilitated by not-for-profit organizations. Vendors also participate and contribute. Vendors benefit when they help develop a successful open standard in this way because their offering is immediately compliant and therefore marketable. A program manager within the DoD POR may or may not choose to consume the vendor’s off-the-shelf product as a system component to comply with the new standard. Meanwhile others in DoD, or anywhere else, are free to use mission funding to procure and sustain the “consumable” product if it suits their mission needs.

The industry component of the NCO Incubator is the W2COG Institute, an incorporated non-government, not-for-profit organization of technically expert providers and consumers of netcentric artifacts from government and industry. The W2COG Institute provides an excellent venue for defining project parameters, selecting and organizing participants, and serving as a clearing house for artifacts, best practices, and lessons learned

The government component of the NCO Incubator should be a “working capital” activity designated as a “software test range.” Such ranges are described in US Title 10 and may accept reimbursable funds from government and industry sources for shared use of range infrastructure. Call this government activity a “unifying netcentric program office” or UN-Program Office (UNPO). The UNPO will facilitate use of DoD laboratories as venues for pilot projects and distribute reimbursable funding accordingly. UNPO activity will include capturing hardware, software, and COI NR-KPP specifications and metrics required to replicate reference implementation results. The UNPO will populate and maintain the Virtual GIG and develop procedures to apply approved NR-KPP specifications and metrics in the VGIG test environment. That process will apply a stamp of approval on demonstrated NET-READY vendor offerings. Further, UNPO will ensure that these NET-READY products are immediately placed on a GSA or other convenient procurement schedule.

The UNPO and the W2COG Institute should be linked with an “Other Transaction Authority” (OTA) or Cooperative Research and Development Agreement (CRADA) contract. The objective of the OTA/CRADA would be mutual discovery and development in the realm of netcentric science. Terms of the OTA/CRADA will require W2COG Institute to provide reimbursable funding through the UNPO for shared use of DoD “Software Range” infrastructure and/or consultation with government technical experts as appropriate. Government activities that wish to charter W2COG Institute pilot

projects will MIPR funds to the UNPO for further transfer via the OTA/CRADA to the W2COG Institute. Neither organization will charge the other “pass through.” *A single contract vehicle will allow low friction, netcentric interaction among multiple government and commercial activities.*

Proposed Incubation Projects

These projects will be designed to rapidly field useful capability within less than a year, and at the same time provide reference implementations of netcentric architectural components that can be used to accelerate the programmatic NCO-enabling engineering process. *The NCO Inc* will broker among operational units, sponsors, government programs and labs, academia, and vendors to put together teams of operators and government/industry engineers to quickly field well-defined solutions to specific operational issues with "shrink wrapped" components designed to be reusable in other applications. *NCO Inc* will carefully select these projects on the basis of the following critical success criteria.

- ☐ Compelling and urgent operational use case whose mission outcome can be demonstrably enhanced by delivery and timely exploitation of more valuable information.
- ☐ Mission performer willing to participate actively on the team.
- ☐ Funding source with "now year" money.
- ☐ Finite, well-defined, deliverable composed of shrink wrapped NCO-enabling components and achievable in < 12 months.
- ☐ Applicable technologies and engineers expert in their use on the team.
- ☐ Viable vendor business model defined to sustain the application and the users.
- ☐ Support by appropriate leadership & legal bases covered.
- ☐ Well-defined dedicated FTE and management structure.

Project participants will be funded to some level, but will be expected to donate some resources as well. Intellectual property donated and developed will be shared, and participants will be protected, according the W2COG IPR policy. Participants will be selected from W2COG members on merit via process sanctioned by W2COG governance policy.

Mobile Collaborative Networks

Where "privacy" can be maintained by participants according to their own situation-based release policies to a level of assurance accepted by, for example, the banking and health industries.

Project 1

Netcentric Humanitarian/Disaster Response

Disaster response and humanitarian aid are critical issues for DoD. The military needs to do better at interacting with multi-national agencies including NGO's, UNOs, and OGOs. Trust is an issue. Communications is an issue. Interoperability is an issue. Net Centric Operations (NCOs) offer some ready made solutions. This need offers an outstanding opportunity to apply existing, and further develop, netcentric capabilities in "Operations Other Than Warfare" (OOTW) – we learn fast when working real world compelling issues!

The BIG IDEA is to use shrinkwrapped COTS communication, collaboration, publish subscribe, and Multi-Level Security MLS) technology to develop plug and play netcentric disaster relief components. These components can then be assembled as required to provide a tailored package for a specific relief or humanitarian mission. We will use W2COG status as a non-threatening NGO to facilitate NGO – military interaction. Initial steps are to provide a demonstration of the capability and environment using the available components, and work with the training being done at California State University, Monterey Bay. As part of the training effort, we will work with the stakeholders to demonstrate an unclassified MLS environment which can be used for mission sharing based on peer-to-peer transaction among mission participants.

Extended Benefit: We will use lessons learned and components developed as NCO architectural reference implementations for other applications. Components include chat tool, privacy technology, pub/sub, modular privacy policy, and router suite.

Approach: Work with vendors, universities, ONOs, OGOs, USGOs and NGOs interested in humanitarian relief. Identify compelling use cases such as training/education and/or e-commerce for disaster survivors or evacuations needs during a government collapse. Create MLS "privacy" domains with COTS packages used by, e.g., commercial entities, governments, and private users. . Include "chat." Include COTS publish/subscribe tool. Load specific client on a lap top for each individual domain participant; provide data and server interfaces as required. Link NGOs and others via "fly away" or wireless router. This will create a mobile collaborative network where "privacy" can be maintained by participants at an extremely high level of assurance according to their own situation-based release policies to a level of assurance already accepted by, for example, the banking and health industries.. Various network components (chat tool, privacy technology, pub/sub, modular privacy policy, and router suite) are all "plug and play" NCO reference implementations that, having been demonstrated to be effective, can be used for other netcentric applications in DoD or elsewhere.

Project 2

"Private" but UNCLAS Cross Coalition Collaboration Network

Terrorists' network is "UNCLAS" and populated by various "open" sources of intelligence. Bad guys' info-sharing is unencumbered by US security policies regarding security risk for information release under need to know. Good guys info sharing ability is encumbered by monolithic security policy associated with classified network. The consistent message from US and coalition participants in Operation IRAQI FREEDOM is that they are at a disadvantage because they are unable to share Critical Time Targeting (CTT) or Putting Ordinance on Emerging Targets (POET).

The BIG IDEA is to use COTS to change bad guys' *asymmetric* info-sharing advantage into a *symmetric* info-sharing advantage to the good guys. Emphasize IED defeat. The key is to use high assurance readily available MLS components to enable a network in which mission critical, but unclassified can be rapidly collected and shared appropriately, according to the need to share requirements of each of the participants.

Extended Benefit: We will use lessons learned and info-sharing components as NCO architectural reference implementations for other applications. Components include: chat tool, privacy technology, modular security policy, pub sub, mobile router

Approach: Work with operational ground unit. Work with NCOIC to identify vendors interested in developing mobile network, cross-domain security, and collaboration products available for commercial high assurance but not necessarily meeting NSA assurance levels for classified information. Work with NASA to leverage their work with mobile routers and multi-spectral sensors. Create multi-level "privacy" domain with COTS packages used by commercial enterprises, such as banks and health organizations, which must comply with many privacy and identity requirements similar, but not at as high an assurance level, as MLS/CDS networks. Include "chat." Include COTS pub/sub. Load specific client on a lap top, provide data and server interfaces as required. Put mobile router and UNCLAS spectral and chemical sensors on UAV, helo, balloon, satellite, etc to establish on demand wireless network and find anomalous signature information generated by hidden IED and generate reports using interoperable semantics. Use NASA and commercial, such as CNN and FOX News, mobile network technology to link sensor image to expert analyst and analyst to operator. Use UNCLAS private cross coalition collaborative network to publish and update IED alerts. This will create a mobile coalition "UNCLAS" network where "privacy" can be maintained by coalition members according to their own situation-based release policies to level of assurance accepted by banking and health industry. Leave classified system as is. Various network components (chat tool, privacy technology, modular security policy, pub sub, mobile router) are all reference implementations that can be used by other operators with similar info sharing requirements.

Valued Information at the Right Time (VIRT)

A critical design objective for netcentric communication architecture is to favor delivery of valued (as defined by the individual user) information at the right time.

Project 3

Valued Information at the Right Time (VIRT)/Rich Semantic Track:

A critical design objective for netcentric communication architecture is to favor delivery of valued (as defined by the individual user) information at the right time. Semantic web technology will be critical to enable the appropriate automated machine-to-machine transactions. DoD's GIG policy addresses these points in its Community of Interest (COI) approach to metadata tagging of "information objects." Information objects, in this context, can include not only traditional stateful objects, such as contact locations; but also more dynamic objects, such as time-rate of change or threshold crossings. There are two aspects of the current approach which create concern. First, many existing notions and definitions of COIs are themselves based on traditional stove-pipe notions and are not well aligned with the intent of NCO. Second, the access to information, such as track or contact information, which is critical for situational awareness, is also based on stove-pipe systems, interfaces, and formats. For example, the "track", i.e. the time rate of change of an object of interest, is a particularly critical information object which is not available as an object, but rather is information in multiple formats and contexts which must be somehow rationalized by the consumer. Further, defining the technical interface specifications for exchanging track and other situational awareness information across the boundary of an enterprise information-sharing domain and closed real-time systems (e.g. weapons or sensors) is critical for GIG success.

The BIG IDEA is to incentivize industry to "productize" semantic web using the DoD "track" use case as the prototype application. Capitalize on the success enjoyed by industry in creating a semantic environment in which common sharing can develop rather than the current method of each COI defining those elements of interoperability in a vacuum, or conversely, at higher levels in the organization, defining them so that are much too complex for implementation. This includes working with rate of change and threshold objects at differing levels of security and developing a common semantic alert approach.k based on participants need to share and resources.

Approach: Work with SIAP, PEO IWS, and SOSCOE (others?) to identify specific key information transactions and the associated "semantic team" expert in the critical information exchange elements. Work with NCOIC to identify vendors interested in venturing IRAD toward semantic web development. Work with National Geospatial-intelligence Agency (NGA) and various geospatial analysis tools and vendors to develop and demonstrate an interoperable semantic web construct for VIRT. Develop a scaleable framework for rich semantic software development, and field a working VIRT prototype for at least one cross domain scenario of interest. Various VIRT components (e.g. ontology model, intelligent agent model, semantic team, track model, pub/sub model) will serve as NCO architectural reference implementations for other applications.

Project 4

Valued Information at the Right Time (VIRT)/Cross Domain Change Detection

Define "change detection" as an automated process to compare newer volumetric data sets with previous versions and alerting operators if and only there is an unexpected change of state. **The BIG IDEA is to implement VIRT by applying change detection to key data bases maintained at various security classifications and at various locations around the world, and then issuing the appropriate tailored alert to the affected operator regardless of that particular operator's bandwidth or security limitations.**

Extended Benefit: We will use lessons learned and VIRT capabilities as NCO architectural reference implementations for other applications. Components include geospatial search tools, visualization tools, collaboration tool, change detection algorithms, pub/sub model.

Approach: Work with National Geospatial-intelligence Agency (NGA). Work with NCOIC to identify vendors interested in developing geospatial analysis, collaboration, and cross domain security products. Link COTS and/or GOTS web-enabled geospatial search tools to an NSA approved virtual machine cross domain solution. Apply COTS geospatial analysis tools to identify user-defined change detection criteria at each level of system security classification. Create VIRT change detection visualization products at each level of system security classification. Publish an unclassified alert message using cross-domain collaborative tool in each case. Various VIRT components (e.g. geospatial search tools, visualization tools, collaboration tool, change detection algorithms, pub/sub model) will serve as NCO architectural reference implementations for other applications.

Distributed, Adaptive Operations

Project 5

Network-Centric Command and Control—A Logistics Application

Effects-based maneuver warfare and distributed, multi-dimensional operations require a high level of operational adaptation. To be effective, the support subsystem must be designed to be an adaptive, network-centric system with dynamic optimization of sourcing options and delivery of critical supplies like ammunition, food, parts, fuel, and even troops all driven by commander's intent. Some companies, e.g. WALMART, have become very successful through optimization across supply chains within their business models. Further, agent based technology can be applied to generate decision support tools associated with supply chain issues. **BIG IDEA: Apply best of breed of commercial supply chain management tools, real-time assessment models, real options theory to integrated supply/service chains (risk management) and agent-based decision support tools (DSTs) to enable distributed, adaptive logistics capability.** Work funded by both DARPA and the Office of Force Transformation has led to important discoveries that can be heavily leveraged toward this end.

Extended Benefit: We will use lessons learned and service oriented mediation technologies as NCO architectural reference implementations for other applications. We will explore how SOAs and agent-based DSTs can increase the speed of command and quality of delivering operational effects. Components include, "rules engine", pub sub, intelligent agent decision support model.

Approach: Work with USMC operational units. Align with NOMADD ACTD. Use technology delivered to OSD OFT Sense and Respond Logistics (SARL) initiative as a base line. Deliver a working suite of network embedded Logistics C2 decision support services.

NETCENTRIC Business Process:

The same concepts that make NCO an effective approach to operational mission accomplishment apply to business process.

Project 6

Netcentric Preparation of Budget Exhibits for multiple budget reviewers:

Because of the large numbers of reviews required for moving budgets through the fiscal process of any large enterprise, there are always a multiple budget exhibits that must be prepared. Although the same basic information is required for each exhibit, the method for presentation and the details of the schedule and the highlights are very often different, depending on the budget reviewer.

BIG IDEA: Leverage the commercial industry investment in fiscal report preparation to enable and facilitate government reporting. Banking and investment firms have a very similar problem insofar as they must report on the same two basic variables, time and money, but in many, many different ways.

Extended Benefit: We will use lessons learned and info-sharing components as NCO architectural reference implementations for other applications. Components include: collaborative tool, privacy technology, modular security policy, re-useable report generation software, pub/sub.

Approach: Work with existing government acquisition and reporting organizations, such as DoD Systems Commands and Program Execution Offices, and with existing commercial banking and investment tools to exploit and deliver a set of resources that can be made available for multiple reporting requirements. In current banking and investment operations, although some of the business rules differ from government requirements, for example regarding taxes, the basic method of exposing the required variables are common. By exploiting this common framework and putting different business logic under the reports, and reusing the methods, we will demonstrate an ability to reduce cost for these reports by reusing existing resources.

15. Managed Service GIG, Public Private Partnership for Netcentric Transformation

MANAGING NETWORK SERVICES

Executive Summary

Achieving the extraordinary improvement DoD seeks by fielding Netcentric Enterprise Services (NCES) will require extraordinary non-traditional adjustments. Service Oriented Architecture (SOA) represents a powerful, but immature C4 paradigm. It has been successfully and incrementally applied in some e-business sectors, but not yet universally. Hence, rather than invent its own solution, DoD should harvest the benefit of SOA technology as it matures following the same approach and realizing the same economies and dexterity enjoyed by the most adept Internet technology consumers, i.e., (a) Capture huge benefit of scale to drive down infrastructure and integration cost; (b) Invest savings massively in innovative business process improvement. Success requires that DoD collaborate in new ways with new partners among the e-business market place's best of breed, leveraging their infrastructure, and fueling their innovative processes in precisely the same way successful firms conduct e-business in myriad sectors like finance, travel, gaming, etc. *The task requires a conceptual mission transformation from builder and operator of networks, to superb enabler and manager of network services* that will give our warriors information advantage over our adversaries. It will require five major change initiatives: (a) Develop a business model and *business targets associated with cost reduction and netcentric productivity*; (b) Structure contractual partnerships as e-biz consumer vice acquisition program manager (govern and incentivize by customer defined objective Quality of Service, security, productivity, and cost metrics vice system specifications); (c) Select prime contractors on the basis of success in e-biz vice defense systems; (d) Develop and govern collaborative V&V methodologies for SOA; (e) Incrementally align process and human resources with the new mission model.

Challenge/Opportunity

NCES is brilliant in that it recognizes both the emergence of a new era in information technology-- the era of *service oriented* distributed computing -- and the need for DoD to embrace it as an a powerful and cost-effective enabler of *netcentric operations*. GIG, and especially NCES, is deliberately intended as a transformational initiative.

“Transformation” means “big change fast.” It is not realistic for any organization, DoD included, to effectively exploit a disruptive technology without undergoing substantial internal disruption. The danger is to get distracted by details of the enabling technologies and lose sight of the underlying objectives. The power of the concept is its notion of guaranteed level of service to multiple disparate consumers from any number of trusted

providers. Therefore the required big change is in focus: concentrate on the guarantee of quality and trustworthiness of the service, not the detail of the technology. We must monitor, invest in, and *govern* increasingly powerful ways to productively share and/or protect information in an increasingly competitive, heterogeneous, and dangerous network landscape. We should not unilaterally build and operate a monolithically “secure” DoD C4 system.

Tom Friedman’s popular book *The World is Flat* describes how savvy organizations have joined in symbiotic Internet-enabled ecosystems to achieve a huge economy of scale and deliver unprecedented levels of success. The DoD objective similarly should be to join appropriate ecosystems and harvest the benefit of technology as it matures following the same approach and realizing the same economies and speed to capability enjoyed by the most adept Internet technology consumers. Conventional wisdom says DoD requirements are just too unique for generic Internet solutions, but conventional wisdom, by definition, is an anathema to transformation.

See TAB B for multiple impressive examples of ultra economies of scale achieved by e-business providers and consumers that allow them to invest massively in business process improvement. DoD organizations can achieve similar economies and opportunity not by competing or copying, but by partnering with the e-business market place’s best of breed, leveraging their infrastructure, and fueling their innovative processes in precisely the same way successful firms conduct e-business in various sectors, including finance, travel, gaming, etc.

The task will require DoD activities to develop a business models and associated business targets, structure the appropriate contractual partnerships, select the right partners, and deal with legacy issues pragmatically. All this must be achieved in bite-sized chunks, and will require creative new approaches to well known problems.

Business Model and Targets

The big change is not about technology *per se*; it is about fielding and using technology more adaptively and collaboratively, deliberately seeking economies of scale, and leveraging emergent strength against emergent opportunity. Auditing and budgeting tools that match these transformational objectives can be implemented rapidly and wholesale. Converting legacy process and human capital to the new world model will require a sustained sense of urgency, and a relentless incremental approach to re-aligning resources.

In one sense, the netcentric task is to *transform into a Great practitioner of e-business* ...the business just happens to be warfighting. Transformation from Good to Great (as explained by Jim Collins in his book of the same name) requires designing a business plan around the intersection of:

1. Your passion
2. Your economic engine
3. The thing you can feasibly do better than anyone else in the world.

Outsourcing an organization's passion is suicide. Outsourcing everything else is fair game. I assume that DoD passion is war fighting success.

The economics of the DoD network are on a scale of tens of \$B/year. The majority of that investment is used to maintain an increasingly archaic, expensive, and *vulnerable* legacy infrastructure. The economic engine of network innovation must be based on the remaining margin. And Peter Drucker would say that the *principal responsibility of top managers is to shift as many resources as possible toward these capability improvements*.

It is unrealistic for the DoD to compete with Google *et. al.* in trying to build the world's best network services. It is entirely realistic for DoD to become the world's best manager of defense-appropriate network services.

Given that DoD C4 overhead is far greater than commercial firms with similar data volume and actual (if not required) quality of security, it is feasible to find resources through improved efficiency. Therefore, the DoD business model should aim to drive down basic C4 infrastructure and integrations costs and re-invest savings in business process innovation, including proactive computer network defense. The first task is to develop a disciplined GIG business model, based on a symbiotic public/private "ecosystem," agnostic about particular technology, which marshals resources and applies them to achieve objectively defined business targets.

ACTION: Draft business plan. Include POA&M for ACTIONS 1-10 described below. (See TAB C)

Objective #1 is to drive down cost of operations and infrastructure. DoD's requirements for communications security and reliability for the majority of its traffic are on par with requirements in specific identifiable industry sectors. E.g., security services that satisfy Sarbanes/Oxley privacy requirements in the financial sector will satisfy a large subset of DoD's less critical security requirements. Drive down costs by embracing open IT standards and outsourcing basic network services via frequent competitive bid. Seek bids from basic Internet service providers (e.g. AOL, Yahoo, AT&T) for their excess capacity, and from value added Internet providers (e.g. Google, AOL, Yahoo, AT&T, Akamai, IA "boutique" vendors, etc.) to customize content *and security services* to DoD requirements.

ACTION #1: Measure, or at least closely estimate, current DoD "network" costs. Objectively quantify the various DoD levels of service and security requirements. Set fiscal targets for infrastructure cost reductions by surveying e-business providers on the basis of cost vs. service delivered.

Objective #2 is to partner with customers and invest cost savings in business process innovation. Again, the approach is to free government human resources to focus on achieving organizational objectives, rather than designing and/or building systems -- manage the forest, don't plant the trees. Incentivize customers to make big changes fast by ensuring they recapture cost savings for re-investment in their critical requirements; make them active champions of this approach vice bystanders.

ACTION #2: Work with customers to set measurable netcentric productivity targets. E.g. a time sensitive strike community of interest (COI) aims to

decrease average time between detection and engagement of targets; a logistics COI aims to decrease its inventory at rest; a training COI aims to improve test scores and decrease time in school; a C2 community aims to decrease decision timeline and increase decision quality.

E-business Contractual Partnerships

Traditionally government program managers carefully design and develop government systems, including C4 systems. Defense vendors build to carefully designed top-down mandated system specifications. Compliance is tested at every milestone. It takes many years to field systems. Equivalent C4 capability is fielded and refreshed much faster in the e-business sectors.

Noticing its speed to capability disparity, DoD has tried in some cases to emulate the e-business model through managed service contracts. This “consumable” acquisition model is much less ponderous than the system acquisition process, and managed service contracts have proven to be effective for life-cycle maintenance of basic capability such as telephone service. However, DoD has had limited success applying the idea to information processing capability because of: (a) Insistence on control of the specific technical solution and (b) Selecting prime contractors who are not expert in e-business transformation. Success on the order of the accomplishments described in TAB (B) will require DoD to ***execute managed service contracts as a true consumer vice program manager.***

A case in point is that the current RFQ for NCES Collaboration Services. The RFQ is much more detailed and constraining than a commercial version would be. It describes use case scenarios but not targeted productivity outcomes. It dictates a solution composed of a good list of standard features and one good price model rather than invite vendor innovation around a problem statement. By comparison, the truly transformational RFQ that TWA issued to build the world’s first airliner was a half page long.

The objective is to perform Enterprise Service Management by designing, monitoring, and enforcing carefully designed high level Service Level Agreements (SLAs) among GIG Service Providers (GSPs) to continuously develop, deploy, and upgrade services via managed service contracts. This model is employed by many of today’s most successful e-businesses (e.g. H&R Block, Lending Tree, Travelocity) and is supported by the kind of “business services” you see IBM pitch in its “On Demand” marketing campaign.

The job becomes to design observable and enforceable SLAs that link directly to military business targets, and then negotiate relatively short term contracts with bonus and extension clauses that are keyed to delivered and measured service. ***DoD shouldn’t care if its contractors employ nine or 109 enterprise services as long we get X.X content delivery, at X.X latency, X.X information assurance at X.X reduced cost; and gets X.X increase in enterprise netcentric productivity units and X.X measured security quality. DoD should care deeply about what those X.X’s should be, especially the ones that describe netcentric productivity, and partner creatively with industry by investing in commercially viable solutions that keep the numbers moving in directions DoD cares about.***

The alternative, DoD attempting to lead the development of web SOA technology by building it like a weapons system, is unrealistic given the DoD's small IT market share and ponderous development process. (That said, government activities should be free to compete with the commercial market place to offer their managed GIG service offerings.)

Picture GIG as a black box that provides GIG services via dynamic knowledge portals. Portals are designed to deliver customer-defined Valued Information at the Right Time (VIRT). Call this idea *my-VIRT*. Customers vary widely. E.g., *my-VIRT* could simultaneously and automatically deliver targeting data to a fast moving jet and its controlling C2 center when a critical pop-up target is identified by an un-attended ground station. Likewise, *my-VIRT* might deliver an IM to a desktop announcing that it's time for the customer's annual dental check up and an opening at the clinic matches a free slot on the customer's calendar. DoD communities of interest (COI), or "ecosystems" as they would be described in e-business jargon, including contracted *innovative service providers*, should partner on the basis of the art-of-the-possible to continuously improve *my-VIRT* for COI purposes. (Note that this *my-VIRT* concept applies to machine-to-machine, machine-to-human, and human-to-human GIG transactions.)

The GIG *innovative service providers* will adapt to specific mission application, but conform to the same kind of customer satisfaction SLAs, and architectural principles (i.e. GIG as an ISP) that DoD dictates, monitors, and enforces. DoD C4 providers should therefore work with their operational customers to define SLAs in terms of specific mission performance improvement targets. *Contractors whose innovative services demonstrably enhance DoD productivity per these mission thread targets should be rewarded with bonuses and renewals.*

ACTION #3: Per objective #1, drive down cost by outsourcing network infrastructure and streamlining network operations. Avoid getting caught up in vendors' technical solution by issuing very short RFQs based on SLAs around price points, reliability, and security for basic GIG services. Include any DoD-owned network infrastructure as government furnished equipment in these contracts.

ACTION #4: Per objective #2 field nces (lower case deliberate) incrementally via innovative services time-and-material clauses in GIG managed service contracts. Use the same legal basis as for life cycle maintenance contracts, i.e. technology refresh. "Maintenance" in the Internet world means "innovatively refreshing technology." Use netcentric productivity targets to design nces SLAs based on improving the value of the information delivered, not the technology that delivers the information.

ACTION #5: Execute NCES Discovery/Mediation/Collaboration pilot project per ACTIONS #5 & #6 above. The VIRT example is not notional. Task the pilot project to leverage the existing VIRT COI (TAB D) The VIRT ecosystem is an open IP environment that includes large and small companies from defense and IT sectors (all of which have contributed IRAD) as well as DoD researchers. The VIRT COI has developed a commercially viable first-of-breed semantic web architecture and is ready to demonstrate a reference implementation. The commercial partners are eager to be first to

market with cutting edge semantic web products and services, which happens to be developed around the DoD requirement.

Choose Contract Partners Carefully

Business logic defined above dictates that DoD select GIG-service providers on the basis of demonstrated success in e-business rather than success in defense industry. Choose prime contractors who understand effective, interoperable, distributed, information processing as e-business providers do. Prime contractors shouldn't be allowed to bid on a piece of the GIG if they can't show quantitative examples of their success, or their sub contractors' success, with respect to the kinds of SLAs DoD requires. The prime contractors who win GIG work shouldn't necessarily be traditional defense contractors.

ACTION #6: Develop selection criteria for GIG service providers.

Collaborative Validation and Verification for SOA

GIG governance requires a fundamental change to C4 system Validation and Verification. Designing, monitoring, refining, and enforcing (through appropriate auditing methods) these SLAs will take hard work and creativity.

This is an opportunity for DoD to take a global leadership role in accelerating the deployment and utility of SOA technology. The fundamental change required is implied by the term "collaborative." By its nature SOA requires "trusted" transactions among "discovered" unknown entities in a dynamic enterprise. *The ability to validate and verify on the basis of risk vs. reward, and need-to-protect vs. need-to-share, criteria in an open, and potentially infiltrated, architecture is crucial, but as yet immature.*

Further, traditional approaches to C4 system V&V stress wringing out all the bugs before *risking* exposure to the customer community. Collaborative Internet approaches to V&V, e.g. as used by the LINUX open source community, stress the opposite approach. That is, *deliberately use the customer community to assist in the V&V process prior to wringing out all the bugs so that on-the-fly adjustments can tweak in favor of greater utility as mutually discovered by partnered customer and provider.* (All the bugs may never be wrung out, but the advantages of a faster rate of evolution more than compensate.)

ACTION #7: Create, govern, and help operate, a global center of excellence for SOA Collaborative Verification & Validation (CV&V) that balances requirements for security, operational continuity, etc., with the need for speed.

Convert Legacy Capability

The best Internet service companies invest massively in innovation, achieving ratios on the order of 10 employees innovating for every employee “keeping the lights on.” If DoD is going to meet the innovative cycle requirement of the GIG, it must achieve a similar ratio. Likewise, the most successful e-businesses have evolved ultra efficient “flat” organizations that enhance netcentric opportunity within their own organizations. DoD Combat Support Agencies that hope to achieve e-business-like flexibility will need to develop similar less hierarchical organizational models.

ACTION #8: Design the ideal work force and work processes for a great netcentric services managing organization. Derive management organization accordingly.

ACTION #9. Set targets for decreased manpower required for basic Computer Network Operations (CNO) and increased manpower for CNO developmental activity

ACTION #10: Develop training and incentive plan to convert current work force and processes into the ideal work force and processes.

C. R. Gunderson

TAB A: Panel of Experts

Geoff Brown, Oracle, geoffrey.brown@oracle.com
Dr Aaron Budgor, Aaron Budgor & Associates, budgora@comcast.net
Mike Daconta, Department of Homeland Security, Michael.Daconta@dhs.gov
John Emerson, Amberpoint, jemerson@amberpoint.com
Mike Friedel, Akamai, mfriedel@akamai.com
Greg Gardner, Oracle, Greg.Gardner@oracle.com
Jack Golden, AT&T, jagolden@att.com
Steve Graves, The Corner Group, steven.graves@thecornergroupp.com
Martin Guttman, Intel Corp, martin.guttman@intel.com
Dr Rick Hayes-Roth, Naval Postgraduate School, fahayesr@nps.edu
Rick Jones, Intel Corp, rick.jones@intel.com
Dawn Meyerreicks, America on Line, Dmyrix@aol.com
David Minton, Planning Systems Incorporated, dminton@plansys.com
John Nagengast, AT&T, nagengast@att.com
Ash Parikh , Raining Data Cor., ash@rainingdata.com
Dr Raymond Paul, ASD NII C2 Policy, Raymond.Paul@osd.mil
Hans Polzer, Lockheed Martin Co, hans.w.polzer@lmco.com
Ajay Ramachandran, Raining Data Corp, ajay@rainingdata.com
Prof Paul Strassmann, Former Director of Defense Information, paul@strassmann.com
Chris Thomas, Intel Corp., chris.s.thomas@intel.com
Erick Von Schweber, Synsyta & Neological, erick@infomaniacs.com,
Linda Von Schwebeber, Synsyta & Neological, linda@synsyta.com
Mike Wolf, Green Hills Software, mwolf@ghs.com

TAB B: Examples of E-business Economy of Scale and Innovation

- ✓ Akamai operates 17,000 servers across 1100 physical networks, in 2500 locations in twenty countries guaranteeing 100% content delivery to 2000 commercial and government firms using a single network operation center manned by six people.
- ✓ AT&T has implemented an order of magnitude greater reliability on its B-t-B service-over-IP than on its voice telephone service.
- ✓ Price points available to AOL and its communication service peers are drastically less than those available to DoD despite comparable market share.
- ✓ Prof Paul Strassmann, former Director of DoD Information, is conducting research on computational infrastructure cost reduction. He has discovered a number of impressive examples including:
 - Google has fielded nearly 300,000 servers in clusters distributed around the world, effectively the largest super computer on the planet, yet the cost of their hardware infrastructure in terms of procurement, installation, maintenance, and operations, is a small percentage of their capital outlay.
 - RightNow Technologies, which operates an extensive global C4 enterprise to execute its Customer Relations Management business model, spends only 6% of its revenue on hosting costs. (per quote from RightNow Technologies CEO Greg Gianforte)
 - Alexa (www.alexa.com a subsidiary of Amazon.com) offers effectively on-demand web-enabled super computer services on an affordable per/use basis. E.g. \$1 per cpu hour ; \$1 per Gigabyte storage per year (with multi-terabyte is available to store user applications, source code and data processing output); \$1 per 50 Gigabyte processing (for data transfers to and from the computers reserved for the processing of applications); \$1 per Gigabyte of data uploaded or downloaded (for data transfers to and from the user's own computers); \$1 per 4,000 web services requests (for using Alexa on-line publication services).

TAB C: List of Actions

ACTION: Draft business plan. Include POA&M for ACTIONS 1-10 described below.

ACTION #1: Measure, or at least closely approximate, current DoD costs. Objectively quantify the various DoD levels of service and security requirements. *Set fiscal targets* for infrastructure cost reductions by surveying e-business providers on the basis of cost vs. service delivered.

ACTION #2: Work with customers to *set netcentric productivity targets* for NCES RFQ's. E.g. a time sensitive strike community of interest (COI) aims to decrease average time between detection and engagement of targets; a logistics COI aims to decrease its inventory at rest; a training COI aims to improve test scores and decrease time in school; a C2 community aims to decrease decision timeline, and increase decision quality.

ACTION #3: Per objective #1, drive down cost by outsourcing network infrastructure and streamlining network operations. Avoid getting caught up in vendors technical solution by issuing very short RFQs based on SLAs around price points, reliability, and security for basic GIG services. Include any DoD-owned network infrastructure as government furnished equipment in these contracts.

ACTION #4: Per objective #2 field nces (lower case deliberate) incrementally via innovative services time-and-material clauses in GIG managed service contracts. Use the same legal basis as for life cycle maintenance contracts, i.e. technology refresh. "Maintenance" in the Internet world means "innovatively refreshing technology." *Use netcentric productivity targets to design nces SLAs based on improving the value of the information delivered, not the technology that delivers the information.*

ACTION #5: Execute NCES Discovery/Mediation/Collaboration pilot project per ACTIONS #4 & #5 above. The VIRT example is not notional. Task pilot project to leverage the existing VIRT COI (TAB C) The VIRT ecosystem is an open IP environment that includes large and small companies from defense and IT sectors (all of which have contributed IRAD) as well as DoD researchers. The VIRT COI has developed a commercially viable first-of-breed semantic web architecture and is ready to demonstrate a reference implementation. The commercial partners are eager to be first to market with cutting edge semantic web products and services, which happens to be developed around the DoD requirement.

ACTION #6: Develop selection criteria for GIG service providers.

ACTION #7: Create, govern, and help operate, a global center of excellence for SOA Collaborative Verification & Validation (CV&V) that balances requirements for security, operational continuity, etc., with the need for speed.

ACTION:#8: Design the ideal work force and work processes for a great netcentric services managing organization. Derive management organization accordingly.

ACTION:#9. Set targets for decreased manpower required for basic Computer Network Operations and increased manpower for CNO developmental activity

ACTION #10: Develop training and incentive plan to convert current work force and processes into the ideal work force and processes.

TAB D: Valuable Information at the Right Time (VIRT) Ecosystem

Member Organizations:

Boeing
The Corner Group
Intel Corporation
L3
Lockheed Martin Company
Naval Postgraduate School
Ohio University
Oracle
Neologic
Pillar Data Systems
Raining Data Corporation
Rockwell Collins
Synsyta
Teknowledge

Summary and status: See attached PowerPoint presentation

16. Collaborative Validation and Verification for Service Oriented Architecture

Collaborative Validation and Verification for Service Oriented Architecture

Today DoD test and evaluation (T&E) process for C4ISR acquisition programs is *system centric*. ACAT programs are driven by formal DoD requirements. DoD specifies performance parameters from those requirements and designs and implements systems accordingly. System developers assemble components, and verify satisfactory incremental and system performance per specific milestones. Prior to fielding it, *operational testing* (OT) authorities install the system on a test platform, and observe as operators verify that it meets system performance specifications under stressful operational loads. OT may be conducted off-line from real operations, but is always conducted in an operational environment.

GIG represents a fundamental change to our C4ISR capability model. GIG is not intended to be a system, but rather a *service oriented* system-of-systems. A simplistic abstraction is that GIG will consist of (a) “plumbing” made of networked hardware, (b) “services” made of software, (c) clients, which may be thought of as hardware and software consumers, such as operators, or providers of services; (d) and middleware, a combination of hardware and software, which links the services and clients together.. Consider software services as independent software modules sitting on independent hardware components. Operators will “discover” these networked software modules. Metadata will describe how to use and combine the service, e.g., may provide service level agreements (SLA). Clients will then combine distributed services as necessary to perform C4ISR operator tasks. “The network”, whether LAN or WAN, becomes the computer. No need to “hardwire” the software modules together inside a box labeled with an acronym like WWMCCS, JOTS, or GCCS.

GIG policy aims to continuously field both plumbing and services to support operators who will also be part of the GIG. Each incremental addition to the GIG should contribute holistically. That is, GIG topology should exponentially expand in ever increasing creative and innovative ways as tools and services are fielded and operational users discover more rapid and more powerful ways to find and use information. *GIG policy specifies that this process will not be controlled centrally and COTS development of the World Wide Web has proved it cannot be centrally controlled successfully.*

This description of SOA CONOP is grossly simplified and leaves out important technical detail. However, the take-away is that *in order to be successful, the GIG must be developed and tested as information-centric, not system-centric*. GIG developers must tackle the technical detail in ways that make this simple vision possible. Therefore, we need an information-centric T&E model to help iron out GIG technology and methodology. Our old C4ISR T&E model aims to assure that unique information processing systems work properly. The new GIG T&E model should assure that networked information is properly treated by increasingly generic networked tools. For example, because much of information sharing is about transactions and domain sharing, the new T&E model should focus on these type of metrics rather than system types of metrics. That is, *we need a T&E model designed to assure that information available on the GIG is (a) protected as necessary, but (b) broadly and increasingly useful.*

The T&E model NSA uses for certification and accreditation (C&A) is instructive in this regard. NSA tests the way information is secured rather than just how systems perform. NSA “type certifies” software associated with a particular Information Assurance (IA)

reference implementation and documents "owner's manual" details associated the certification particulars, for example configuration options and interfaces. A program manager can get an NSA security certification without invoking the rigorous unique system requirement by, (a) using the same "type" off-the-shelf IA software, and (b) adhering to the NSA "owner's manual". With this certification in hand, it is then easier to get an operational accreditation for a specific installation or implementation.

The new GIG T&E model can leverage the NSA C&A model over a much broader spectrum of information processing capabilities and services. The "net-enabling" C&A process should field reference implementations to develop and validate netcentric specifications, and then verify that a candidate software "type" process, resource, or capability satisfies them. For example: does the software comply with GIG information-centric technical policy; does it align with COTS standards, trends, and investments; are there transition risks; is it useful in mission context (e.g. power requirements, user friendliness, bandwidth reality, durability, size/weight); does it measurably enhance netcentric information and transaction productivity?

This new GIG *net-enabling* C&A process should co-evolve with NSA's Information Assurance (IA), or, *net-protecting*, C&A process. The two together, *net-protecting* C&A plus *net-enabling* C&A, will constitute *Net-Ready* C&A. U.S. Title 10 has a provision that authorizes DoD to operate software test ranges, and accept funding from commercial companies for their use. *Net-ready* C&A process should exploit that provision and rigorously apply the distributed suite of DoD network test facilities to evaluate candidate commercial and government reference implementations. The evaluation should generate "owner's manuals" that document how to install and to use the "type certified" net-enabling software to achieve the validated and verified netcentric benefit.

A program manager who bundles components that are type-certified as *net-ready*, and satisfies an audit on the test range that the bundling is consistent with the approved owner's manuals, will earn a net-ready accreditation for his system. The remaining question is how to perform information-centric OT of this *system* accredited as *net-ready*? As discussed, the GIG will be a service oriented system-of-systems. Therefore, the OT model should consider a candidate "system" as a GIG service -- a *service-of-services*. This "system-of- services" abstraction scales infinitely.

DoD should design tests to verify that the new GIG system/service adds value to the enterprise per the targeted netcentric productivity specifications. That kind of verification can only be performed collaboratively with other GIG services, and the customer community. Fortunately, because the new service has already been accredited as *net-ready*, the OT can be conducted "on-line." Successful e-business practitioners use this kind of on-line collaborative T&E. They *invite their customers to help wring bugs and make on-the-fly improvements*. They have learned that all the bugs may never be wrung out, and that operational customers have limited time to spare to assist in development, but that the advantages of a faster rate of evolution are sufficiently compelling. DoD can capitalize on their lesson by following their model per the following proposal.

NET-READY CERTIFICATION OFFICE (NCO)

Proposal: Establish a government working capital activity, a Net-ready Certification Office (NCO), which will serve as a distributed test range for Certification and Accreditation of network-enabling off-the-shelf software per the following characteristics:

1. Coordinates activity of existing distributed assets rather than create new infrastructure or bureaucracy.
2. Provides government activities immediate access to broad spectrum of industry, government, and academic information processing experts and products via standing contractual vehicle(s) with not-for-profit W2COG Institute
3. Facilitates rapid formation of “ecosystems” of operators, vendors, labs, and sponsors formed around compelling operational issues, mature technology, and ~90 day spirals to demonstrate, prototype, and productize bundled information processing software.
4. Assists operational units and other customers to develop netcentric productivity metrics, i.e. measurable increase in the value of information available to the operators as a result of a new or “improved” tool.
5. Coordinates authorizing, scheduling, and funding for commercial use of DoD network test facilities.
6. Engages appropriate un-biased warranted authorities to test according to the following criteria:
 - a. Compliance with published SOA/GIG technical reference and policy documents.
 - b. Alignment with COTS standards, trends, and investment
 - c. Transition risk assessment including schedule
 - d. Utility in mission context (e.g. power requirements, user friendliness, bandwidth reality, durability, size/weight)
 - e. Netcentric productivity
7. Certifies and/or accredits that off-the-shelf network-enabling software satisfies government requirements including documentation of reference implementation details: hardware, software, and interface specifications; training and doctrine requirements.
8. Facilitates placing certified and/or accredited software on approved consumable procurement schedules.
9. Establishes and maintains a “GIG-lite” distributed on-line repository of off-the-shelf network enabling software (both certified and pending certification) and mission thread based modeling and simulation demonstration suite.

10. Maintain open source library of contributed code generated in NCO activities.

Requirements:

Staff:

1. Project Manager. (1 FTE X 2.5 years. Six months research + Two years operational implementation = 2.5 years. Full time government employee.)
2. Chief Engineer (1 FTE X 2.5 years. Government employee or contractor. Dedicated or shared time.)
3. Software range manager (1 FTE X 2.5 years. Government employee or contractor. Dedicated or shared time.)
4. Contracts officer (1/2 FTE X 2.5 years. Government employee. Shared time.)
5. Marketing officer (1/2 FTE X 2.5 years. Government employee or contractor. Dedicated or shared time.)
6. Administrative support (1/2 FTE X 2.5 years. Government employee or contractor. Dedicated or shared time.)

Governance:

1. OSD mandate
2. Working capital designation

17. Hard Problems in Network Operations

Addressing the Hard Problems in NetOps: A White Paper Posing Novel Solutions to Overcome Tradition-Based Thinking

Preface: NetOps, the three-component approach to modern communications support to military operations using the Global Information Grid (GIG), is in the midst of what might be considered its “sophomore year.” GIG Network Management, GIG Network Defense and GIG Content Management (up until recently known as Content Staging) comprise the three core components of NetOps.⁹⁴ NetOps promises to improve military communications by assured system and network availability, assured information protection and assured information delivery.⁹⁵ Although the concepts and components are essentially a given at this point, this White Paper seeks to ask strategic questions about NetOps and the GIG and thereby help us pass the entrance exams into our “junior year” and beyond. We must begin to exploit what we have been exploring to this point. We must apply what we have learned, and execute graduate-level NetOps as soon as possible.

Author Jeff Cares recently asked in his book, Distributed Network Operations,⁹⁶ whether or not we are “getting it right” with Net Centricity and Net Centric Operations, major conceptual justifications for the GIG. Considering our nation’s investment in the GIG and how hard we’ve worked to push the power of networked communications to the edge of the GIG, Jeff’s is a timely and strategically significant question. Let us then frame the hard NetOps problems with two fundamental categories of questions: how do we know we are getting it right (i.e., doing the right kinds of things); and, given that we are doing the right things, what can we do to effectively and continuously improve on our investment and support to our nation’s warfighters? In summary, then, the framing questions for considering NetOps Hard Problems become: “are we doing the right things and are we doing them right?”

Are we doing the right things?

Fundamentally, the first question we must ask is do we think in terms of network centricity or do we think in terms of knowledge centricity?⁹⁷ Is it about the network or is about the information that flows within the network? Why does the network exist in the first place? Is the warfighter interested in how the network does its job, or rather if the network does the right job (i.e. provide especially useful information)? The answers to these basic questions revolve around the old management saw: *do the right job, then do it right*; effectiveness first, then efficiency. Frankly, in the midst of our sophomore year, it is not yet clear then that GIG NetOps is about doing the right job first and foremost. We’ve gone to some length to seek efficiencies in what we think is “good NetOps,” but we must challenge ourselves by asking if we are network focused or knowledge focused. Most of our technological innovations in support of NetOps suggest that we are network focused. The network is necessary, but not sufficient to provide graduate-level support to the warfighter.

⁹⁴ Relevant acronyms used throughout this paper: GIG Network Management – GNM; GIG Network Defense – GND; GIG Content Management – GCM.

⁹⁵ JTF-GNO Draft Global information Grid NetOps ConOps, Version 3, with comments, March 2006.

⁹⁶ Cares, Jeff, Distributed Network Operations: The Foundations of Network Centric Warfare, 2005, Alidade Press, Newport, RI

⁹⁷ Although one could argue the distinctions between data, information and knowledge, those distinctions are not necessary here – this is not a knowledge management paper. The terms data, information and knowledge are used interchangeably in this paper.

On the other hand, relevant and timely knowledge, whether delivered by the GIG or not, is both necessary and sufficient to support the warfighter. By now, we are compelled to believe the GIG is necessary to deliver knowledge, but let's not confuse efficiency with effectiveness. It is useful to consider the hypothesis that GIG Content Management, the least defined and least mature component of NetOps, is actually the most important of the three. Network Defense exists not only to protect the network, but also to protect the information that flows within it. Network Operations exists to ensure the network is available and efficiently managed, again to provide an access path for the information the warfighter needs. If there was no information available, would the warfighter invest time and resources in the GIG? The answer is "not likely." As Lt Gen Charles Croom, commander of the JTF-GNO has said in many public forums, "it's about the information, not the network." It follows that GIG Content Management, or at least a knowledge-centric focus, is the right thing to do.

According to the most recent approved version of the GIG NetOps ConOps, Global Content Management (called Information Dissemination Management /Content Staging in the current version) is the "...technology, processes and policy necessary to provide awareness of..." relevant and accurate information available to users of the GIG. Its core services include content discovery, content delivery and content storage.⁹⁸ On behalf of the commander, these core services seek to:

- Permit commanders to adjust information delivery methods and priorities for enhanced Situational Awareness
- Allow information producers to advertise, publish and distribute information to the warfighter
- Enable users to define and set information needs (profiles) to facilitate timely and efficient information delivery and/or search information databases to retrieve desired products as required
- Improve bandwidth use
- Enhance all aspects of the GIG transport capabilities⁹⁹

Having established that the network should in fact play second fiddle to the information it transports and protects, we must still confirm the criticality of the network and the assets it can bring to bear to do the right job. However, our new recognition of the most important elements of effectiveness, should change the way we prioritize our efforts and marshal resources. For that reason alone, it is critical to make certain we are doing the right job in the first place. GIG Network Management and GIG Network Defense are therefore two important elements of doing the job right, but let us acknowledge that they soak up the bulk of resources *to the exclusion of doing the most important job* at the dawn of GIG-based net-centric operations.

In order to ensure we graduate to our junior year, we need to properly align our priorities. Fortunately, many of the general concepts that support GIG Content Management also support operations and defense. We can use what we already know and have available.

⁹⁸ JTF-GNO Joint Concept of Operations for Global Information Grid NetOps, Version 2, 10 August 2005 (Final Approved Version).

⁹⁹ *Ibid.*

Hence, NetOps is the right thing to do on behalf of warfighters to guarantee access to knowledge they need. Now let's look at what we can accomplish to "do the thing right."

What can we do to improve?

To align effort and investment across DoD, we must define "improvement" in universally understood and measurable terms. Such a definition requires clear delineation of the productivity we hope to achieve through knowledge-centric NetOps. Oddly, this definition of productivity does not yet exist. However, if delivery of relevant and timely knowledge is the necessary and sufficient condition for success, then it follows that productivity will increase as the value of the bits exchanged increases, where "value" is specifically defined in terms of relevance and timeliness for particular applications. The principle applies to both offensive and defensive aspects of what we could call "agile NetOps." An example of this agility, or metaphoric "maneuver," would be demonstrated if we could measurably subtract value from the bits exchanged by adversaries through diversionary or disruptive techniques, achieving the same effect as adding value to our own exchanges. We will discuss maneuver and mobility/counter-mobility as part of the "Other Questions," below.

Almost all NetOps and Information Assurance conferences raise the issue of metrics. Effective collection and assessment of metrics could provide the baseline for significant progress that we have been missing in NetOps all along. Even the claim that situational awareness of the GIG is the "holy grail"¹⁰⁰ of NetOps moves beyond hyperbole with the addition of metrics. Meaningful metrics help us know what we are actually visualizing and what difference modification and adaptation really make. One of the first steps we must take to improve NetOps and solve the "doing it right" framing problem is to introduce NetOps Instrumentation and Mensuration. Instrumentation is a straightforward concept, even if we don't do it well yet. Mensuration, on the other hand, a term borrowed from the intelligence and mathematics communities, is where the even harder work starts.

To set the stage for how to instrument the GIG, and more appropriately GCM, let's discuss what should be measured. The term "mensuration" comes from the studies of geometry and calculus, denoting the measurement of areas, volumes and the lengths of curves.¹⁰¹ The Imagery Analysis community uses this term to indicate precise measurement of volumetric surfaces (e.g., multi-dimensional) when obtained from an apparent flat surface such as a photograph. Mensuration also seems to accurately capture the idea of precision measurement of network geometries such as the GIG expresses.¹⁰² "...mensuration, in its practical aspect, is of importance for giving reality to the formulae themselves and to the principles on which they are based..."¹⁰³ Translation to "reality" is the exact problem we face in measuring the success or failure of the GIG. The GIG is a high-dimensional network that actually defies even the term "grid" and therefore measurement requires techniques that transcend simple matrix or lattice measures of axes. The term "mensuration" seems to align well with this challenge.

¹⁰⁰ Observations from both the JTF-GNO J3 and J5, in various briefings and conferences since the inception of the NetOps CONOPS, June, 2004.

¹⁰¹ "Mensuration," *LoveToKnow 1911 Online Encyclopedia*. © 2003, 2004, LoveToKnow. Found at: <http://25.1911encyclopedia.org/M/ME/MENSURATION.htm>, accessed 29 March 2006.

¹⁰² For an excellent primer on complex network geometry, see Cares, pp. 149-172.

¹⁰³ *Ibid*, "LovetoKnow"

But “value” is the operative term regardless of the specific techniques we use to assess that value. We just need to understand value in a consistent manner that translates across the community. Accordingly, we need to design NetOps instrumentation, modeling, simulation, and mensuration techniques as part of the process to analyze the parameters of relevance and timeliness that define value for the various NetOps communities of practice (both friend and foe!).¹⁰⁴

This NetOps process analysis will provide a holistic systems engineering perspective on the myriad programmatic elements of the GIG, a perspective that will highlight “the most important jobs” and the associated critical system components. The “most important jobs” in NetOps will define the vital functions in the supporting acquisition process from basic research through life cycle maintenance. It will also map the critical and supporting interactions that produce the rich complexity of emergence we seek in this approach.¹⁰⁵

An example of complexity-producing interactions reflects a convergence of disciplines such as those that compose NetOps. The complexity of mixing GND and GNM produced at the same time more challenges for the JTF but offered deeper insight into effective operations and defense of the GIG. Like Clausewitz’s wrestling analogy where two wrestlers can achieve conditions together that they cannot achieve apart, the fusion of GNM and GND allowed the JTF to achieve new insights that were not possible before.¹⁰⁶

Now add the third component of NetOps, GCM, to this convergence and we have the potential for even greater richness and opportunities for warfighter support that could never have happened if the components were kept apart. And, as we have argued throughout this paper, GCM offers to provide the greatest opportunity for enhanced support to warfighters and their supporters throughout the entire enterprise. At the very least, we should be able to substantiate this claim once we have discovered the sweet spots of interaction and emergence in content management as a function of NetOps...and, once we can affix value to these interactions as manifested by content “points of presence” within the warfighter’s domain.

Success in measuring value in the basic NetOps functions of GCM, GNM and GND can cascade throughout all GIG-related functions. Armed with that information we thus can design, execute, and iterate a DoD knowledge-based business plan that will balance investments in network management, defense, and content most productively.

Evidence-based value?

The GIG, and most importantly GIG Content Management, must be accurately measured to ensure we are doing the job right. What should we measure in the content management component to ensure we are doing the right job? There are several conventional

¹⁰⁴ There is considerable value in considering the enemy as part of the Community of Interest (or ecology) of the GIG, a point worthy of additional research, as well.

¹⁰⁵ In this paper, “complexity” is a positive trait rather than something to be avoided. It reflects interactions and emergence as functions of sophisticated organizations and processes that demonstrate the capacity to evolve and generate adaptive capabilities. See Cares, 2005, and Kauffman, 1995. There is a rich body of research and application in the area of complex adaptive systems which this paper and the accompanying proposal will seek to leverage.

¹⁰⁶ Bassford, C., “Clausewitz and His Works,”

<http://www.clausewitz.com/CWZHOME/CWZSUMM/CWORKHOL.htm>, accessed 19 April 2006.

measurements that contribute to understanding the value of content to warfighters and those who support them. These measures include statistics on number of times accessed, number of times edited or forwarded, utility voting metrics (e.g., “popularity votes” in the commercial internet environment),¹⁰⁷ and value calculated from “willingness to pay” for certain information (a type of economic-based schema that may also be useful in the GCM world). Other measures that are less conventional but require examination to determine their potential in assessing value include:

- productivity enhancements related to GCM and the other NetOps components; a further discussion on productivity measurement follows below
- entropy-related measures that reflect declining or expiring utility of knowledge or increased disorganization of that knowledge and its components
- organizational- or mission-pattern matching measures that infer other likely customers of knowledge to predict where information should be staged in the GIG
- economic-related measures that reflect trade, upgrade, or other value-added functions concerning the construction of new and inferred knowledge¹⁰⁸
- semantic web-related ontologies and taxonomies that also predict relationships of definitions and context that will assist in staging knowledge throughout the GIG¹⁰⁹

Warfighter productivity must increase in measurable ways or the GIG and NetOps are not only inefficient, they are ineffective. An early action we must take now is to analyze, measure and report upon the productivity the GIG has delivered and offers to warfighters in the near- and long-term. The warfighter’s access to, and use of, knowledge provide us especially important mensuration points of GIG effectiveness, and thus explains why this paper stresses GCM as the key parameter to measure. Productivity, as reflected in the attached proposal becomes the initial thrust for instrumentation and mensuration of NetOps and the GIG.

Instrumentation for mensuration of the non-linear and irregular surfaces of the GIG require novel approaches to tool development. Jeff Cares suggests what he calls “The Information Age Combat Model” as a way of understanding the complex interactions of weapons and sensor platforms in a non-linear combat model. Tools that evolve out of Cares’ networks of nodes will also affect mensuration of NetOps. The initial network nodes he suggests include the following:

- sensors: objects that receive and transmit signals about observed phenomena
- deciders: objects that receive data from sensors and determine the present and future arrangement of other nodes within the network

¹⁰⁷ The use of “live votes” on www.msnbc.com is an example of this approach. Although hardly scientific, this method allows readers to judge in coarse terms the value of an article or predict some outcome in a very quick poll. This information has problems in terms of accurately assessing true value, but can provide top-level insight when other metrics are not available or too difficult to attain. When taken as part of a broader profile, these metrics may become more valuable as we become accustomed to using them.

¹⁰⁸ Inferred knowledge refers to an emerging intelligence community concept known as evidence-based inference. See Andrews, Twining and Schum, 2005. David Schum has published numerous papers on the use of evidence organization techniques to support the generation of new evidence and knowledge. These techniques offer great promise for GIG Content Staging and will be explored in detail in the project proposed in this white paper. An automated version of this technique was proposed in “Agent Based Evidence Marshaling”, in Hunt, 2001.

¹⁰⁹ Future proposal presentations will deal in significant detail about how value may be calculated as part of multiple interactive metrics and as independent metrics.

- influencers: objects that take direction from deciders, interact with other nodes and take some action that affects the states of those other nodes
- targets: objects that are nodes on the network that are not sensors, deciders or influencers¹¹⁰

Cares' "Information Age Combat Model" suggests an entirely new family of measurement devices that can sense and transmit information about the actions taking place on the GIG within all three components (GCM, GND and GNM). These tools will also enable eventual machine-to-machine interaction to stimulate self-healing and self-synchronization, two important objectives of the next generation Global Information Grid, according to Lt Gen Croom. These mensuration points also offer insights in measuring productivity relative to GCM, as well.

NetOps is not a uniquely military concept. On the contrary, there are many practitioners of virtually identical ideas in the e-business sector. The successful Internet Service Providers (ISP) all invest massively in innovative process improvements to gain increasingly rich knowledge of network state in order to ship more and better information content, and to protect against both denial of service and unintended disclosure. Clearly there are superb existing industrial capabilities that the DoD can simply port or outsource. Arguably, since all major ISP's contract and/or partner with a company called Akamai to help them perform "InternetOps," Akamai with its "God's eye view of the Internet" represents best of breed in this regard.¹¹¹ Akamai stages information content at the tactical edge of the Internet in servers strategically located on the basis of traffic analysis. This approach prevents the need for Internet consumers to access Akamai customers' core infrastructure. Further, Akamai monitors activity and seamlessly routes packets in the most efficient and effective manner and employs a suite of productivity metrics to continuously improve the service. By using its distributed sensor technology to perform post-time diagnostic audits of malicious network activity, Akamai can help clients perform agile network "counter maneuvers". Akamai is in 2,500 locations, more than 1,100 networks, in 70 countries and routes 20-25% of all traffic on the World Wide Web. Their NOCC is manned by only six people whose principle task is to deploy and monitor new applications "live" on the network.¹¹²

So, what can we do better? According to Professor Chris Gunderson, Naval Postgraduate School, "In order to tackle the technical challenges described in earlier paragraphs, we should certainly consider following the industrial e-business model of deliberately driving down infrastructure costs, and then reinvest massively in innovative business process. To do that, we must re-think our current tightly controlled contractual model in favor of more enlightened partnership and incentivized managed service contracts that can harness the IT industry's innovative baseline to address government requirements on Internet time scales."¹¹³

¹¹⁰ Cares, pp. 75-106, with a description of a littoral battle scenario on pp. 110- 122.

¹¹¹ *Wired Magazine*, issue 11.07 Jul 2003, <http://www.wired.com/wired/archive/11.07/slammer.html>, accessed 5 April 2006

¹¹² Quoted from Mike Friedel, Akamai Director Sales, Public Sector, by Professor Chris Gunderson, Naval Post-Graduate School, 5 April 06.

¹¹³ Quoted from Professor Chris Gunderson, Naval Post-Graduate School, Executive Director of the World Wide Consortium for the Grid, 6 April 2006.

Other Important Questions Related to Hard NetOps Problems:

Initial discussion of hard problems generally poses more questions than it answers. Accordingly, to resolve the hardest of the NetOps problems, we must ask questions that help us understand and model the complex process interactions around those problems. Only then can the problems be solved by invoking non-intuitive, multi-discipline approaches. A representative sampling of these questions includes the following. Questions without answers, or at least hints, are of limited use to busy people. Hence, simple beginnings to possible answers accompany the questions.

- What are the key interactions and linkages that must occur to ensure the GIG is doing the right job? Doing the job right?

Until this point in history, the DoD has tended to treat GCM, GND and GNM as separate entities. Only with the advent of the JTF-GNO did we introduce an entity theoretically designed to integrate all three constructs. While we have made limited progress integrating GND and GNM, there is still much to do. We have only just begun to consider how to integrate the third element, GCM. Further, all progress to date has been halting and discrete. Meaningful success will require a new effort, perhaps the new PEO-NetOps within DISA, to pull all three components together so that progress in one area deliberately contributes positively, and by design, to the other areas. The vision will be complete when all three components are systemically enhanced by improved processes that support each other in a continually synergistic feedback loop. We will observe key interactions and linkages through deliberate NetOps planning and design models and simulations that require all components to complement each other.

- Are the current blend of GNM, GND and GCM the right components for NetOps today and tomorrow?

As a conceptual framework, it's enough for now. Certainly we need to apply emerging technologies such as semantic web capabilities (ontologies and service-oriented architectures) effectively to make the concepts increasingly useful for warfighters and their support teams. A future component candidate for NetOps might include Mission Mapping as a separate but fully integrated capability that GNM, GND and GCM all affect equally well.

- What value does the JTF-GNO add to NetOps?

To build what the JTF-GNO J5 calls NetOps Forces we must explore this question in detail. The JTF-GNO does add limited value today, but when mixed with the right blend of relevant forces, it can emerge as a full-partner warfighter in GWOT and all other operations that our Combatant Commanders will undertake in the future. Effective compositions of NetOps Forces are key to success of the GIG and NetOps, and deserve treatment within a separate paper that could be delivered in future treatments of the NetOps Hard Problems. Examples like Akamai, as discussed above, offer clues on how NetOps forces may eventually emerge. Today, however, the JTF-GNO adds value as an organizing entity, as a first generation capability to integrate the current components of NetOps, and as a research lab for warfighters to assess the effects of a NetOps-enabled Global Information Grid.

- Is Net Force Maneuver a useful construct?

If information is the commodity of interest; if information dominance is the objective of the game; if “the network” is the playing field; and if our adversary is clever, then it follows that we need a dynamic and multi-faceted strategy to win both the battle and the war. The idea behind maneuver is to pose a fundamental shift in thinking about defense and operation of the GIG, as well as suggest the themes of mobility and counter-mobility for all three current components (GCM, GND and GNM). There are emerging technologies that have already been tested for worm and virus defenses, as well as diversionary tactics, that offer substance to this new form of network warfare. However, we must refine the questions about maneuverability as well as test the answers to these questions in ways disciplined by NetOps thinking that focuses all three components in synergistic ways.

- What new acquisition techniques can we develop to facilitate a rapidly evolving GIG that exploits disequilibrium rather than fall victim to it? I.e., how do we collaboratively build the GIG so that useful co-evolution occurs in manageable ways to dampen the effects of disruptive technologies and attacks?

Wake up one day and check out what’s happening at JTF GNO. Then go to the Internet and take a look at all the top of the line personal information processing equipment you can purchase on line: cell phones, TiVOs, iPods, laptops, ISP services, Internet sites, Travel & Financial services, etc. Go to sleep and wake up a year later. Check out what’s happening at JTF GNO. Then go back to the Internet..... There’s an answer in there somewhere; it has a lot to do with the way we come to define “manageable.” Our acquisition model must be relevant and co-evolve with the warfighter’s needs. The Internet thrives in a constant disequilibrium state that naturally imposes its own limitations and dampening effects, without particularly strict rules and design criteria apart from data and transmission standards. There are clearly examples resident in the emergence of the World Wide Web and the Internet that offer us insights in how the GIG will emerge in future iterations. We will likely witness Government-academic-commercial collaborations materialize that will streamline acquisition policies and techniques. We must build environments that accommodate the development of simple rules that more closely mimic natural “acquisition” systems to ensure co-evolution occurs in meaningful ways that work on behalf of the warfighter and the underlying support system.

- What are the appropriate risk mitigation techniques that we can use to reduce the apparent fragility of the GIG and how do we measure success of these techniques if they work and prevent disaster from occurring in the first place?

Once again, modeling and simulation will be of great value to help us visualize the threats and potential mitigations we can bring to bear in risk management and overcoming the fragile nature of the GIG. We have built the GIG and it’s linkage to warfighters using the best technologies available, but we must better understand the dependencies that have grown up around our communications networks (or any network for that matter). And, to avoid experimentation with the real network and warfighter’s bread-and-butter communications systems, we must rely on disaster and recovery testing, as well as threat prevention *in silico*. From these models, we will develop tactics, techniques and procedures to deal with problems before they occur. We must design

these models so that they reveal linkages and dependencies and provide early indicators and warnings that reflect the real world.

Summary

This white paper proposes that there are hard problems for the NetOps community of interest to solve in order to maximize early, if limited, success in implementing the first years of NetOps. We have framed these problems in the context of doing the right thing and then ensuring we are doing it right. We are essentially sophomores in NetOps years and have learned enough to make us either dangerous or appear that we know more than we really do. We must now ask the very hard questions about our purpose and intent, acting on behalf of the warfighter, and harness a blend of art and science within a sophisticated modeling and simulation environment. As an extension to asking hard questions, we must also measure key parameters such as enhanced productivity to ensure we are doing the right job in the first place. This will reveal both shortcomings and strategies in ways that cost less money but produce worthwhile insights. It is also likely that success in our consideration of effectiveness and efficiency in NetOps will reveal meaningful ways to pose the requirements necessary to build the next generation GIG. The accompanying proposal demonstrates a viable way-ahead in accomplishing this purpose.

Prepared by: COL Carl Hunt, Ph.D., US Army, Director of Technology, JTF-GNO.

Assisted and reviewed by Professor Chris Gunderson (CAPT, USN, Ret), Naval Post Graduate School, Executive Director of the World Wide Consortium for the Grid.

18. Command Resiliency: An Adaptive Response Strategy for complex Incidents

Command resiliency : an adaptive response strategy for complex incidents / Joseph W. Pfeifer

Pfeifer, Joseph W.

Electronic access: <http://library.nps.navy.mil/uhtbin/hyperion/05Sep%5FPfeifer.pdf> (377 KB)

Personal author: [Pfeifer, Joseph W.](#)

Title: [Command resiliency : an adaptive response strategy for complex incidents / Joseph W. Pfeifer.](#)

Publication info: Monterey, Calif. : Naval Postgraduate School ; Springfield, Va. : Available from National Technical Information Service, 2005.

Physical description: xvi , 71 p. : col. ill. ; 28 cm.

General note: Thesis Advisor(s): Christopher Bellavita.

Dissertation note: Thesis (M.A. in Security Studies (Homeland Security and Defense))--Naval Postgraduate School, September 2005.

Bibliography note: Includes bibliographical references (p. 67-70).

Abstract: Many organizations believe they are prepared for the next terrorist event by wrongly assuming there is a predictable threat that can be managed with the purchase of new equipment. Unless organizations develop a resilient response strategy that can adapt organizational and operational elements to respond to new terrorist incidents, they will find themselves with the same difficulties emergency responders did on 9/11. As terrorist attacks unfold, organizations are pushed beyond their normal capabilities. How quickly organizations adapt to the uncertainty of a new crisis is critical. Organizations that cannot adapt to new threats of large, complex terrorist events will be less likely to respond effectively to future attacks. This paper recommends a resilient response strategy that is flexible enough to adapt to complex incidents. It proposes policy recommendations that address organizational strategy and operational crisis management to deal with the initial critical hours of a terrorist attack. Organizational strategy defines core competencies and what happens when competencies are pushed beyond their capacity. Operational crisis management will examine situational awareness requirements, flexible decision-making and innovation. Command resiliency is achieved by overcoming organizational bias and integrating organizational preparedness and operational adaptability into a synergistic response network.

Technical details: System requirements: Adobe Acrobat reader.

Local note: Fire Department City of New York author (civilian).

Subject: [SUBJECT](#)

Corporate author: [Naval Postgraduate School \(U.S.\)](#)

Other forms: Also available online.

19. Hastily Formed Networks

Hastily Formed Networks

The ability to form multi-organizational networks rapidly is crucial to humanitarian aid, disaster relief, and large urgent projects. Designing and implementing the network's conversation space is the central challenge.

On Sept. 11, 2001, terrorists attacked the World Trade Center, taking 2,749 lives. The attack resulted in severe economic impact, especially to airlines, and a stock market loss of \$1.2 trillion. On Dec. 26, 2004, a tsunami from a 9.1 earthquake overran the shores of many countries along the vast rim of the Indian Ocean. Over 283,000 people died. On Aug. 29, 2005, Katrina, a category-5 hurricane, knocked out electric and communication infrastructure over 90,000 square miles of Louisiana and Mississippi and displaced 1.5 million people. Six months later, New Orleans still housed fewer than 100,000 of its original 1.2 million residents. On Oct. 8, 2005, a magnitude-7.6 earthquake devastated the Kashmir region of Pakistan, killing over 87,000 people. Besides being unexpected major disasters, these events had one other common feature: they all involved hastily formed networks that quickly mobilized, organized, and coordinated massive humanitarian responses.

ROBERT NEUBECKER

The severity of these disasters drove home an important point: the quality of the response depended not on response planning or on new equipment, but on the quality of the network that came together to provide relief. How quickly were voice and data communications restored? How well did the many players from disparate organizations collaborate? How effectively

did the network deliver help to the victims? These incidents demonstrated sharp differences in the quality of the hastily formed network (HFN), which directly affected the effectiveness of the response. Noting that these networks almost always involve military, civilian government, and non-government organizations, the U.S. Departments of Defense and Homeland Security have

The Profession of IT

made it a priority to learn how to effectively assemble HFNs. We coined the term at the Naval Postgraduate School in 2004.

The lessons learned from the networks involving government carry directly into private settings. They will benefit any urgent network of multiple organizations with no common authority that must cooperate and collaborate.

Hastily formed networks is an area where advanced networking technology and human organization issues meet. They can work well together, or they can clash. Our purpose here is to give an overview of this critical area and the challenges it offers to computing professionals.

ORIGINS

The idea of quickly forming a team for a particular, urgent task, and then disbanding it when done, is not new. Table 1 lists three categories of events for which an HFN must respond. Because it involves relatively small teams and known networks, the first category is the easiest and least likely to stress the HFN.

The middle category is the type that emergency agencies such as police and fire departments prepare for. They have professional, highly trained teams ready to respond to particular incidents. They have well-developed practices for advance planning, training in appropriate skills, and positioning of equipment. They already use terms like "ad hoc network" and "crisis response network" to describe what they do.

Category	Characteristics	Examples
K: Known	Know what to do Use existing network structures May choose not to respond	Fast response team for time-critical business problem or opportunity
KU: Known Unknown	Know what to do Don't know time or place Responding network structure known	Local fire, small earthquake, civil unrest, military campaigns
UU: Unknown Unknown	Don't know what to do Don't know time or place Responding network structure unknown	9/11 attack, other terrorist attacks, large earthquakes, major natural disasters (Note: KU events can become UU events when scaled up to large areas or populations)

Table 1. Kinds of events requiring response from hastily formed networks.

The third category puts the greatest stress on the HFN. These events require response beyond the control and capabilities of any single agency. The network structure will depend on the event and the responding organizations.

The main aspects of the third-category challenge are:

- **Genuine surprise.** The precipitating event is in no known category. There has been no advance planning, training, or positioning of equipment.
- **Chaos.** Everyone is overwhelmed. No one understands the situation or knows what to do. People are frantic and panicky.
- **Totally insufficient resources.** Available resources and training are overwhelmed by the magnitude of the event.
- **Multi-agency response.** Several agencies must cooperate in the response, including military, civilian government, and private organizations. These groups have had little or no prior reason to collaborate. The shock of moving from a state of "coexistence" to a state of "collaboration" can be overwhelming.
- **Distributed response.** The response is distributed over a geographical area into many local

jurisdictions. The authority to allocate resources and reach decisions is distributed among many organizations. Decisions by command-and-control do not work.

- **Lack of infrastructure.** Critical infrastructures such as communications, electricity, and water do not work. Makeshift infrastructures must be deployed quickly.

HFN DEFINED

The first priority after the precipitating event is for the responders to communicate. They want to pool their knowledge and interpretations of the situation, understand what resources are available, assess options, plan responses, decide, commit, act, and coordinate. Without communication, none of these things happens: the responders cannot respond. Thus the heart of the network is the communication system they use and the ways they interact within it. We call this the "conversation space" of the HFN.

An HFN has five elements: it is (1) a network of people established rapidly (2) from different communities, (3) working together in a

shared conversation space (4) in which they plan, commit to, and execute actions, to (5) fulfill a large, urgent mission.

An HFN is thus much more than a set of organizations using advanced networking technology. To be effective in action, HFN participants must be skilled at:

- Setting up mobile communication and sensor systems;
- Conducting interagency operations, sometimes called “civil-military boundary”;
- Collaborating on action plans and coordinating their execution;
- Improvising; and
- Leading a social network, where communication and decision making are decentralized, and there is no hierarchical chain of command or ex officio leader.

Most participants do not have a need for these skills in their individual organizations. When they come together, therefore, they find it difficult to accomplish these tasks. When combined

with the overwhelming nature of the urgent event, these inherent difficulties can lead to a breakdown in the conversation space.

CONVERSATION SPACE

The ongoing need to communicate and coordinate is fundamental for the success of any HFN. The term conversation space was introduced for the medium in which all this takes place—from forming community responses to delivering actions. The conversation space is (1) a medium of communication among (2) a set of players (3) who have agreed on a set of interaction rules. These three aspects are summarized in Table 2.

One of our early conclusions was that the effectiveness of the HFN rests on the quality of the conversation space established at the outset. It is not a foregone conclusion that an effective HFN can be established even when the players are trained professionals, as the situations in New York City after 9/11 and in New Orleans after Hurricane Katrina illustrate. In New York, the mayor understood intuitively that

success would depend on everyone, including especially the residents of the city, feeling included in the relief effort. He made sure information was shared, even if piecemeal. While there were some initial coordination difficulties, the network came together and was effective in relief and recovery. A different picture occurred in New Orleans. The various agencies had major difficulties in coordinating and the Federal Emergency Relief Agency (FEMA) did not deliver what people thought it had promised. At all levels there was a lot of finger-pointing and wrangling over who would do what and who would pay for what. When the president put a new man in charge at FEMA, there was no immediate improvement in effectiveness or criticism of the agency. Attempts to impose standard military-style command-and-control in Louisiana and Mississippi were ineffective. This is not intended as a criticism of New York, Louisiana, or Mississippi officials, but rather an illustration that effective coordination may not happen even when all the parties want it to happen.

Certainly a major difference between New York and New Orleans was the sheer scale of the event. New York lost infrastructure in a limited area of perhaps 100 square blocks. The primary agencies in the network ultimately reported to the mayor. Police and fire radios provided basic communications in the “ground zero” area. In contrast,

Table 2. Components of conversation space.

Category	Characteristics	Examples
Physical systems	Media and mechanisms by which people communicate, share information, and allocate resources	Telephone, power, roads, meeting places, supplies, distribution systems
Players	Players included and their roles, core competencies, and authorities	Citizens, fire department, police department, highways department, federal emergency management agency
Interaction practices	Rules of the “game” followed by the players to organize their cooperation and achieve their outcomes	Situational awareness, sharing information, planning, reaching decisions, coordination, unified command and control, authority, public relations. (Note: environment has no common authorities, no hierarchy, many autonomous agents, decentralized communications)

The Profession of IT

New Orleans lost an entire city and was part of a large area (90,000 square miles) with severely damaged infrastructure. All communication systems were knocked out; and as they were gradually being restored, the limited-bandwidth channels were overwhelmed by sheer numbers of citizens trying to use them. Many more agencies had to cooperate on the response. Coping with all this effectively was completely outside most responders' experience.

New York City quickly built trust among the responders and citizens. New Orleans experienced considerable difficulty in building trust.

But this is one of the lessons: the more overwhelming the event, the more likely turf-asserting tendencies will occur and interfere with the effectiveness of the network.

The overarching lesson is: the effectiveness of an HFN depends as much on the participating people and organizations as it does on the communication system through which they interact.

CONDITIONED TENDENCIES

It is well known that individuals under severe stress often forget their recent training and regress

to old, ingrained habits [1, 7]. Richard Strozzi Heckler calls these old habits conditioned tendencies [6]. The old habit is likely to be inappropriate for the current situation and to make matters worse.

The National Institute of Standards and Technology ([4], p. 174) concluded that "a preponderance of evidence indicates that emergency responder lives were likely lost at the World Trade Center resulting from the lack of timely information-sharing..." Police radio transcripts cited by NIST indicate that NYPD helicopters monitoring the two burning towers detected signs of structural collapse in the North Tower and issued an emergency evacuation order to all police. Yet no one in the police department communicated the imminent-collapse information to the fire department. What accounts for this bizarre behavior?

Joseph Pfeifer, a deputy assistant chief in the New York City Fire Department, gives in his master's thesis a detailed example of conditioned tendencies instilled by emergency-response organizations, which paradoxically can render them incapable of effective response in an emergency [5]. Pfeifer was among

those responding to the 9/11 disaster in the World Trade Center. His explanation for non-communicative behavior was that organizational biases—ingrained social habits of the separate organizations—prevented emergency personnel from talking to one another. One of these biases is organizational social identity that prefers to share information within the group but not outside. Under stress, the group members do not think to collaborate or share information outside the group, or to take personal responsibility for the welfare of members of other groups.

The purpose of Pfeifer's study was not to assign blame for needless loss of life in the 9/11 attacks, but to recognize the organizational conditioned tendency as a real phenomenon that can disable an HFN. The question is how to prepare organizations to work together in an HFN and avoid the conditioned tendency. Pfeifer proposed that the agencies use unified command networks, in which leadership is shared among different organizations; for example, an executive committee. This practice will likely create the foundation for HFNs that do not suffer non-communication leadership paralysis.

The effectiveness of an HFN depends as much on the participating people and organizations as it does on the communication system through which they interact.

A GUIDE TO EFFECTIVE HFNS

(1) The quality of the conversation space is critical to success. The space includes the communication systems, the participants, and their interactions within these systems. Effectiveness in conversation space rests on skills that participants may not ordinarily learn in their separate organizations.

(2) The physical communication systems are part of conversation space. Plan and test mobile technologies that can be set up quickly when the regular infrastructure is down. Arrange for security forces to protect the temporary infrastructure. Use and test all communications equipment regularly. Use standard software and protocols—interoperability and simplicity of interconnection will be important. Web services are a good example.

(3) The participating organizations are another part of conversation space. Each brings its own culture, standard practices, and decision-making protocols—which may be incompatible with other organizations. Individuals can become disoriented when familiar organizational practices are suspended. They fail to take initiative, while waiting for orders that will never come. They do not know how to function when there is no common authority, their established command-and-control practices do not work, and collaboration, not control, is the only way to get actions done.

(4) Information glut will be a problem in the network. As com-

munications are initially restored, the victims will overload the severely limited bandwidth as they try to communicate with their families. The responders themselves will overwhelm their colleagues with situational reports and other data. New technologies will be needed to manage information glut and keep the network functioning.

(5) Understand and practice the effective technologies for collaborative networks. These include Web servers to distribute information, wiki and discussion-thread software, chat and instant-messaging services, virtual markets, and coordination services such as Groove (but Groove is restricted to Windows platforms).

(6) Prepare to overcome the barriers to interorganizational collaboration. These include conflicting missions, unclear roles, turf protection, incompatible processes and information systems, disparate cultures, accountability, mistrust, and lack of knowledge of others' capabilities [3].

(7) Prepare for organizational conditioned tendencies to appear under overwhelming stress. Train group members in the basic HFN skills. Promote political support for the organizations to cooperate, mutual respect for the competencies that each organization brings, concern for each other's welfare, and personal responsibility for actions and outcomes. Practice with "unified command"—an executive committee representing the participating organizations that respects the

core competencies that each organization brings.

(8) Train the skill of improvisation. This is a challenge for normal rule-oriented agencies.

CHALLENGES FOR COMPUTING PROFESSIONALS

HFNs bring new words and concepts such as conversation space, coordination without hierarchy, and conditioned tendency. Learn these concepts; they are important.

Interoperability and simplicity are key technical challenges for HFNs. Services offered via Web interfaces are highly interoperable; anyone can use them from any computer. Chat and text messaging services are highly interoperable. But many key services are not. For example, many responders have found the Groove software to be useful for coordination, but Groove runs only on Windows computers; responders with Sun workstations, Apple Macintoshes, or Linux-based computers are out of luck. Many wireless networks are not fully interoperable, for example, Linux and Apple machines use different protocols from Windows machines for encryption and passwords.

To prevent information glut, we need tools to model, label, and filter information so that network participants receive the information most likely to be valuable to them. In crises especially, the participants need to make most effective use of the limited resources of decision-making time and communica-

tions bandwidth by restricting the flow of unimportant bits.

Hayes-Roth shows how a 100,000-fold drop in information volume is easily achievable without loss of effectiveness [2].

Learn about the organizational issues, such as collaborative cooperation and managing conditioned tendencies. You may be part of an organization that responds in an HFN, and you will need to know this.

Understanding how to create HFNs is one of the most challenging parts of modern networking. It is about how a network, its people and its equipment, may function efficiently under extreme stress. ■


REFERENCES

1. Barthol, R.P. and Ku, N.D. Regression under stress to first learned behavior. *Journal of Abnormal and Social Psychology* 59 (July 1959), 134–136.
2. Hayes-Roth, F. Two theories of process design for information superiority: Smart pull vs. smart push. *Command and Control Research and Technology Symposium: The State of the Art and the State of the Practice*. San Diego, CA, U.S. Department of Defense, Command and Control Research Program (CCRP) (2006, to appear); www.nps.edu/cebrowski/Docs/06reports/CI-06-001.pdf.
3. Hovevar, S., Jansen, E., and Thomas, G. *Building Collaborative Capacity for Homeland Security*. NPS Report (Dec. 2004); www.nps.navy.mil/Research/04techrpt.html.
4. National Institute of Standards and Technology. *Federal Building and Fire Safety Investigation of the World Trade Center Disaster: The Emergency Response Operation*. Washington, D.C. (2005).
5. Pfeifer, J. *Command Resiliency: An Adaptive Response Strategy For Complex Incidents*. Naval Postgraduate School. MS thesis, Sept. 2005; library.nps.navy.mil/uhtbin/hyperion/05Sep%5FPfeifer.pdf.
6. Strozzi Heckler, R. *Anatomy of Change*. North Atlantic Books (1984, 1993).
7. Weick, K. *Sensemaking in Organizations*. Sage Publications (1995), 102.

PETER J. DENNING (pjd@nps.edu) is the director of the Cebrowski Institute for information and innovation and superiority at the Naval Postgraduate School in Monterey, CA, and is a past president of ACM.

© 2006 ACM 0001-0782/06/0400 \$5.00

20. Brochure from W2COG Symposium



**World Wide Consortium
for the Grid**

The Inaugural W2COG Working Symposium

GIG Innovation:
Aligning Operator,
Engineer, and Program
Manager Perspectives



24-26 May 2005

George Mason University
Prince William Campus
Manassas, Virginia

The Inaugural W2COG Working Symposium on Global Information Grid Innovation

To Register: www.reconline.com/eventinfo.asp?EventId=22145

The Inaugural W2COG Working Symposium for Global Information Grid Innovation

An Invitation from Dr. Peter Denning, Conference Chair

My Fellow NCO Professional:

Just as the needs for "connectivity" drove the construction of the Internet in the 1980s and "information sharing" drove creation of the web in the 1990s, so is "network centric operations" (NCO) driving the formation of the "Global Information Grid" (GIG) in the 2000s. NCO means organizations employing increasingly distributed operations to perform; the GIG provides the advanced communications network necessary to sustain NCO's highly distributed operations.

The GIG presents a host of technical challenges because NCO in a defense and homeland security environment requires new functions from the network infrastructure -- notably security, bandwidth guarantees, intermittent channels, and collaboration support.

But the biggest challenge of all is social -- involving users, service providers, vendors, manufacturers, and acquisition personnel in deciding what the GIG can offer and how current networks can evolve into the GIG. To meet the technical and social challenges, we must accomplish three things.

First, we must rally all affected communities -- user, technical, and acquisitions personnel -- into a broad coalition of professionals, who share technology, talent, and time, in mutual dedication to advancing NCO.

Second, we must create an open, transparent process for capturing and integrating the work of NCO experts and fostering NCO consensus and innovation. We seek a forum in which government, industry, academia, and international allies can collaborate in forming recommendations on how to deal with NCO technical and acquisitions issues.

Third, we must augment our forum with online and offline resources that simplify the actions organizations must take to integrate NCO technologies into their own operations. These and other resources will be dedicated to nurturing successful NCO adoption.

Our Mission Starts Right Now!

The World Wide Consortium for the Grid (W2COG) is stepping up to provide this forum. Founded in 2004, the W2COG is supported by government, commercial, and academic institutions focused on realizing the benefits for NCO through active resource sharing.

The W2COG convenes its first Working Symposium between 24-26 May 2005. This is not a sit-and-listen-to-presentations symposium, but a collaborative, working meeting to explore NCO engineering issues with peers. It is the ground floor on which we'll build the commitments, processes, and teams constituting the W2COG's NCO mission.

Join the W2COG team, fellow NCO professionals, and me at George Mason University's Prince William Campus to initiate the W2COG's NCO journey.

Dr. Peter Denning
Chairman of the Naval Postgraduate School Computer Science Department
Dir. of the Cebrowski Institute for Information Innovation and Superiority

About the W2COG

The World Wide Consortium for the Grid (W2COG) is a forum for sharing global information grid (GIG) expertise and technology to speed the transformation to network centric operations (NCO). We're comprised of leading government, industry, and academic groups that share the belief that great innovation springs from open, urgent collaboration, independent of any central authority. Our Working Symposium mission is simple: initiate the W2COG's promise for accelerating NCO innovation in the operator, engineering, and acquisitions communities through educational demonstrations of GIG elements, facilitated interaction among NCO professionals, and establishment of open engineering groups required to execute the W2COG's work.

Web: www.w2cog.org

Email: info@w2cog.org

Phone: 703-262-5332

The Inaugural W2COG Working Symposium on Global Information Grid Innovation

To Register: www.regonline.com/eventinfo.asp?EventId=22145

Our mission is universal.
Wherever the competitive arena, winning demands distributed, network centric operations.
Our mission is urgent.
The adversaries we face are more subtle, more dangerous, and more unpredictable.
Our mission starts right now!
Great innovation springs from open, urgent collaboration.

A New Symposium for Meeting New Challenges

A Total Focus on Executing Change

The W2COG Working Symposium will catalyze new relationships, concepts, technologies, and services that can be directly applied to realizing the advantages of NCO and implementing the GIG.

Morning "Innovation Experience" Presentations

Each morning, NCO leaders will present their real experiences driving NCO and GIG change in the operator, engineer, and program manager communities. These sessions will divulge NCO-related strategies and tactics that are working to create mission and campaign advantages today.

Afternoon "NCO Simulation Workshop" Sessions

Each afternoon, the W2COG will provide an example GIG infrastructure and selected mission threads, thereby simulating realistic NCO situations. During these *NCO Simulation Workshop* sessions, NCO professionals from the field will demonstrate their offerings before attendees and a panel of technical experts to reveal essential design, planning, implementation, and adoption experience.

Who Should Attend

- Chief Information Officers
- Chief Technology/Engineering Officers
- Chief Knowledge/Learning Officers
- Civilian and Military Operations and Communications Officers
- Program Managers
- System Acquisition Leads
- Consultant Practice Leads

Benefits of Attending

- Inject ideas into the emerging GIG.
- Tap the emerging stream of NCO innovation by drawing from GIG demonstrations.
- Collaborate with fellow NCO professionals who can help meet NCO challenges.
- Converse with DoD's GIG leaders about requirements, directions, and timelines.
- Interpret NCO directions with leading GIG policy and program experts.
- Design and pioneer a new approach to fielding information technology innovations in the DoD – and beyond.

A True Working Symposium

The W2COG Working Symposium will be a unique opportunity to work directly with emerging GIG solutions – meaning live data, live networks, and live applications – and the professionals that field them. This event is not designed for listeners, but doers: the professionals responsible for executing missions, building implementation teams, fielding solutions, and driving change.

Real Peer-to-Peer

Our Working Symposium objective is to give attendees an opportunity to explore the technical and social challenges they'll face as they champion NCO concepts and GIG technologies. Additionally, the W2COG will utilize these sessions to begin establishing both the missions and make-up of its working groups.

Sponsors



Other Sponsors

- Naval Postgraduate School
- Association for Enterprise Integration
- ASD/NII
- OFT
- DUSD

In Cooperation With

- ACM

Location Details

George Mason University
Prince William Campus
Verizon Auditorium
Occoquan Building
10900 University Blvd.
Manassas, VA 20110
United States

The Inaugural W2COG Working Symposium on Global Information Grid Innovation
To Register: www.regonline.com/eventinfo.asp?EventId=22145

W2COG Inaugural Working Symposium for GIG Innovation 24-26 May 2005 Agenda			
	Day One: The Operator's Mission	Day Two: The Engineer's Challenge	Day Three: The Program Manager's Duty
7:30am-8:00am	*Registration* <6:30am-8:00am> Continental Breakfast	European Breakfast	Continental Breakfast
8:00am-8:30am	Keynote — NCO: The Next Phase of Computing in Defense <i>Dr. Peter Denning</i>	Previous Day Reflections <i>Dr. Barry Frew</i>	Previous Day Reflections <i>Mr. Peter Burris</i>
8:30am-9:10am	Changing Practices Under Conditions of Extreme Stress <i>Maj. Carl Oros</i>	Global Information Grid (GIG) Applications <i>Mr. Geoff Brown</i>	Program Managers as Squires: Getting Today's Knights Their Shining Armor <i>VADM (Ret) Denny McGinn</i>
9:10am-9:50am	The Complex Information Environment of Emergency Response and Recovery <i>Dr. Peter Friedland</i>	Extreme Deployment for Truly Extreme Circumstances <i>MGEN Ricky Lynch</i>	Stories from Way Behind the Lines, on the Acquisitions Front <i>VADM (Ret) Jerry Tuttle</i>
9:50am-10:20am	Collaboration Break	Collaboration Break	Collaboration Break
10:20am-11:00am	Network-Centric Operations: One War Fighter's Perspective <i>RADM (Ret) Winston Copeland</i>	A Brain-Like Computer for Cognitive Applications: The Ersatz Brain Project <i>Dr. Jim Anderson</i>	Integrated Supply Chains: Taking Friction Out of the Acquisitions Process <i>Mr. Mike Daconta</i>
11:00am-11:50am	Innovation Panel Discussion and Summary <i>Mr. Gary Haycox, Mr. John Santini</i>	Innovation Panel Discussion and Summary <i>Dr. Jens Pohl, RADM (Ret) Winston Copeland</i>	Innovation Panel Discussion and Summary <i>Mr. Phil Charles, Mr. Satnam Alag</i>
11:50am-1:00pm	Lunch and Collaboration	Lunch and Collaboration	Lunch and Collaboration
1:00pm-2:15pm	NCO Simulation Workshops: Field Intelligence <i>3 Simultaneous Demonstrations Each Day</i>		
2:15pm-2:30pm	Collaboration Break	Collaboration Break	Collaboration Break
2:30pm-3:45pm	NCO Simulation Workshops: Humanitarian Disaster Relief <i>3 Simultaneous Demonstrations Each Day</i>		
3:45pm-4:00pm	Collaboration Break	Collaboration Break	Collaboration Break
4:00pm-5:15pm	NCO Simulation Workshops: Border Control <i>3 Simultaneous Demonstrations Each Day</i>		
5:15pm-5:30pm	In Recognition of VADM (Ret) Arthur Cebrowski	Open	Collaboration Break
5:30pm-6:00pm			Closing and Plenary Session
6:00pm-8:00pm	Evening Social		Open

W2COG Working Symposium Presentations and Workshops

Day One: The Operator's Mission



Conference Keynote – NCO: The Next Phase of Computing in Defense

Dr. Peter Denning, Chair of NPS CS Department and Dir. of Cebrowski Institute, NPS

Just as the needs for "connectivity" drove the construction of the Internet in the 1980s and "information sharing" drove creation of the web in the 1990s, so is "network centric operations" (NCO) driving the formation of the "Global Information Grid" (GIG) in the 2000s. NCO means organizations employing increasingly distributed operations to perform; the GIG provides the advanced communications network necessary to sustain NCO's highly distributed operations. In this keynote, Dr. Denning will consider both the technical and social challenges facing the community of operators, engineers, and acquirers tasked to make NCO and the GIG a reality.

Changing Practices Under Conditions of Extreme Stress

Maj. Carl "Deiter" Oros, Faculty, Naval Postgraduate School

The types of conflicts facing the U.S. war fighter are changing rapidly. Forces must be lighter, more mobile, and better prepared to act. This goes way beyond planning. Fighters must be intimately connected to information resources, regardless of location or source, critical to sustaining battle context and tactical advantages. This session will review the broad objectives of NCO, laying out priorities and timelines from an operator perspective.

In Recognition of VADM (Ret) Arthur Cebrowski

Vice Admiral USN (Ret) Arthur Cebrowski is considered to be the "Father" of modern theory of Net Centric Operations and Warfare. In gratitude for his service, the NCOIC, the W2COG, and the Association for Enterprise Integration (AFEI) will pay tribute to him on the evening of Day One.

The Complex Information Environment of Emergency Response and Recovery

Dr. Peter Friedland, Asst. Director for Technology & Chief Technologist, NASA Ames Research Center

Instantaneous access to knowledge from a wide variety of sources is key to effective response to and recovery from major disasters of any form, whether caused by natural or human forces. Typical situations involve a dozen or more agencies, commonly with non-interoperable communications and data systems. This talk will survey the information access and management issues from both the field and the operations command center perspectives.

Network-Centric Operations: One War Fighter's Perspective

RADM (Ret) Winston Copeland, Corporate Account Manager, Sun Microsystems

Combat for the war fighter continues to be all about stress. Much has changed over the last 100 years, but historically our proportional losses by friendly fire have not been significantly altered between the Industrial Age and the Information Age. Let's examine Von Clausewitz's Principles of War; are they obsolete? Finally, a perspective on the Top 10 Combat Lessons Observed.

Innovation Panel Discussion and Summary

Facilitator: Dr. Barry Frew, President and CEO, Frew & Associates

Innovation Panel: Mr. Gary Haycox, Director of Strategic Initiatives, Intel;

Mr. John Santini, Chief Scientist, Anteon

The W2COG believes that great innovation springs from open, urgent collaboration. Facilitated by a senior member of the W2COG executive team, each Innovation Panel will spur open collaboration among attendees, highlighting key messages from the morning sessions, and setting the stage for the afternoon *NCO Simulation Workshops*.

NCO Simulation Workshops: Field Intelligence

**3 demonstrations
will simultaneously
run each day from
1:00pm-2:15pm**

The Objective: Provide cross-coalition access to vital intelligence by a warrior on the ground.

Scenario: A Joint Special Operations Task Force (JSOTF) is operating in a hostile environment supported by a military intelligence team. The intelligence team is tasked to continuously obtain and evaluate sensor intelligence data, assess threat and opportunity, and share results with staff and operational units as appropriate. 2 X Predator and 1 X JSTARS Unmanned Aerial Vehicles are shared theater assets whose services can be requested by JSOTF. The team communicates over a 50KBPS SECRET-high circuit.

Sample Value Proposition: Increase the lethality of the kill chain by breaking down the administrative barriers separating warriors from advantageous information, regardless of classification or who holds it.

The Inaugural W2COG Working Symposium on Global Information Grid Innovation

To Register: www.reconline.com/eventinfo.asp?EventId=22145

W2COG Working Symposium Presentations and Workshops

Day Two: The Engineer's Challenge



Previous Day Reflection

Dr. Barry Frew, President and CEO, Frew & Associates

Global Information Grid Applications

Mr. Geoff Brown, Director, GRID Technologies (TBU Advanced Technology Solutions), Oracle

Crafting NCO capabilities for today's war fighters requires engineers to develop an array of new skills, and to let go of a few traditions. Many global information grid technologies (GIG) technologies are disruptive, yet must be mastered and exploited without compromising operator effectiveness. This session will explore GIG concepts, technologies, and capabilities from an engineer's perspective, highlighting the requirement to utilize distributed development to rapidly field distributed systems.

Extreme Deployment for Truly Extreme Circumstances (Teleconference)

MGEN Ricky Lynch, Allied Joint Force Command Naples, Deputy Chief of Staff—Operations, United States Army

General Lynch will speak very frankly about where we are and where we need to be regarding effective information sharing with our allies and coalition members. From an operational, engineering and acquisition perspective General Lynch will discuss recent lessons learned along with addressing critical issues and delineating part of the problem.

A Brain-Like Computer for Cognitive Applications: The Ersatz Brain Project

Dr. Jim Anderson, Professor, Dept. of Cognitive and Linguistic Sciences, Brown University

The Ersatz Brain Project wants to develop a preliminary hardware design, programming techniques, and software applications for a brain-like computing system. The design is based on the "network of networks" approximation to mammalian neocortex which assumes the basic computing unit in cortex is not a single neuron but groups of neurons operating together to form attractor neural networks, roughly at the scale of a cortical column.

Innovation Panel Discussion and Summary

Facilitator: Mr. Peter Burris, CIO, W2COG

Innovation Panel: Dr. Jens Pohl, Executive Director, Collaborative Agent Design Research Center; RADM (Ret) Winston Copeland, Corporate Account Manager, Sun Microsystems

The W2COG believes that great innovation springs from open, urgent collaboration. Facilitated by a senior member of the W2COG executive team, each Innovation Panel will spur open collaboration among attendees, highlighting key messages from the morning sessions, and setting the stage for the afternoon NCO Simulation Workshops.

NCO Simulation Workshops: Humanitarian Disaster Relief

3 demonstrations will simultaneously run each day from 2:30pm-3:45pm

The Objective: Execute 'sense and respond logistics' and 'command and control' in support of third world disaster relief.

Scenario: After a large-scale natural disaster in SE Asia, a humanitarian effort is undertaken to provide relief and stability in a devastated and remote region of a mountainous country. The disaster has eliminated roads and airfields used for accessing the backcountry. The government of the country has requested aid, and will permit US military forces into its borders to assist with initial relief efforts. Additionally, non-government relief organizations are rallying to the cause and are being permitted to enter the country.

Sample Value Proposition: Increase the speed of support to chaotic zones by employing intelligent agents against rapidly accumulating raw data to accelerate evaluation of potential courses of action.

The Inaugural W2COG Working Symposium on Global Information Grid Innovation

To Register: www.reconline.com/eventinfo.asp?EventId=22145

W2COG Working Symposium Presentations and Workshops

Day Three: The Program Manager's Duty



Previous Days Reflection

Mr. Peter Burris, CIO, W2COG

Program Managers as Squires: Getting Today's Knights Their Shining Armor

VADM (Ret) Denny McGinn, Vice President for Strategic Planning, Battelle

Admiral McGinn relates his experiences with the Navy's floating battle lab, humanitarian disaster relief and transformation within the Pentagon to strategic imperatives associated with making information technology work for mankind.

Stories from Way Behind the Lines, on the Acquisitions Front

VADM (Ret) Jerry Tuttle, President, JOT Enterprises

The "Father of JOTS" developed and fielded the first common operating picture by focusing on the operational imperative and breaking paradigms as required. He will offer some thoughts on the way ahead for C4I acquisition.

Integrated Supply Chains: Taking Friction out of the Acquisitions Process

Mr. Mike Daconta, Metadata Program Manager, Department of Homeland Security

The author of [The Semantic Web](#) will share insights about how semantic technology applies to the challenge of fielding NCO capability.

Innovation Panel Discussion and Summary

Facilitator: Dr. Rick Hayes-Roth, Information Sciences Professor, Naval Postgraduate School

Innovation Panel: Mr. Phil Charles, Technical Director, SPAWAR Systems Center, Charleston;

Mr. Satnam Alag, Chief Architect, Rearden Commerce

The W2COG believes that great innovation springs from open, urgent collaboration. Facilitated by a senior member of the W2COG executive team, each Innovation Panel will spur open collaboration among attendees, highlighting key messages from the morning sessions, and setting the stage for the afternoon *NCO Simulation Workshops*.

NCO Simulation Workshops: Border Control

3 demonstrations will simultaneously run each day from 4:00pm - 5:15pm


Objective: Establish international border control.

Scenario: The international intelligence community reports that Al Qaeda "Chatter" is high. It is the height of European tourist season and the Euro is very strong. Airports are thronged. The European Union, the United States, and most of the members of the United Nations have agreed to collaborate with respect to sharing data that might help identify terrorists at border check points.

Sample Value Proposition: Prevent terrorist movement by cross referencing distributed biometrics, stolen documentation, wanted persons, data bases, etc. in real time and in alignment with the myriad international agreements governing behavior.

The Inaugural W2COG Working Symposium on Global Information Grid Innovation

To Register: www.reconline.com/eventinfo.asp?EventId=22145



World Wide Consortium for the Grid

The Inaugural W2COG Working Symposium on Global Information Grid Innovation 24-26 May 2005

Invited Speakers

- Dr. Jim Anderson
- Geoff Brown
- Phil Charles
- RADM (Ret) Winston Copeland
- Mike Daconta
- Dr. Peter Friedland
- Satnam Alag
- Gary Haycox
- MGEN Ricky Lynch
- VADM (Ret) Denny McGinn
- Gabe Minton
- Maj. Carl Oros
- Dr. Jens Pohl
- Mr. John Santini
- VADM (Ret) Jerry Tuttle

Price and Registration Information

Registrations for the W2COG Working Symposium will be accepted only online. To register, please follow these steps:

- **Step 1:** Go To www.regonline.com/eventinfo.asp?EventId=22145
- **Step 2:** Pay \$695US using one of these credit cards



- **Step 3:** Book your own travel and accommodations

For additional information, please call (703) 262-5332.

Location Details

George Mason University
Prince William Campus
Verizon Auditorium
Occoquan Building
10900 University Boulevard
Manassas, VA 20110
United States

Local Accommodations

The following accommodations are provided for informational purposes only. The W2COG cannot facilitate reservations. There is a block of rooms held until 3 May 2005 at the Courtyard Marriott.

Hampton Inn Manassas	Courtyard Manassas Battlefield Park
7295 Williamson Boulevard	10701 Battlevue Parkway
Manassas, VA 20109	Manassas, VA 20109
Phone: 703-369-1100	Phone: 703-335-1300
www.hampton.hilton.com	www.marriott.com/courtyard

Cancellations, Substitutions, and Requests for Refunds

All cancellations, substitutions, and requests for refunds must be completed in writing. Substitutions are welcome. If requesting a substitution, payment already received will be transferred to cover the substitute attendee.

For a full refund, less a \$75 administrative fee, registrants unable to attend the symposium must email their cancellation or refund request to the attention of Michelle J. Belcher at michelle.belcher@w2cog.org prior to 1 May 2005. For cancellations or refund requests made between 1 May 2005 to 15 May 2005, a refund of 50% — less the \$75 administrative fee — will be granted.

A \$75 administrative fee will be applied to ALL cancellations. No refunds will be issued for cancellations received after May 15, 2005, 5:00 PM EST. Refunds will not be issued for no-shows.

The Inaugural W2COG Working Symposium on Global Information Grid Innovation

To Register: www.regonline.com/eventinfo.asp?EventId=22145

Initial Distribution List

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, CA
3. Mr Terry Pudas
Office of Force Transformation
Department of Defense
1401 Wilson Blvd., Ste 301
Arlington, VA 22209
4. Major General Charles Croom
Director, Defense Information Systems Agency
701 South Courthouse Road
Arlington, VA 22204-2199
5. Mr Richard Lee
Deputy Undersecretary for Advanced Systems and Concepts
Defense 3700 Pentagon
Pentagon Room 3D285
Washington, DC 20301-3700
6. Mr Linton Wells
Assistant Secretary of Defense of Network Information Integration
6000 Defense Pentagon
Washington, DC 20301-6000
7. Dr Jonathan Smith
Defense Advanced Research Project Office
3701 Fairfax Dr
Arlington, VA 22203
8. Dr. Peter Denning
Computer Science
Naval Postgraduate School
Monterey, CA 93943
9. Christopher Gunderson
c/o CNRI
1895 Preston White Dr
Reston VA 20191

10. Sue Higgins
Cebrowski Institute
Naval Postgraduate School
Monterey, CA 93943
11. Dr Mark Pullen
Director, C4I Center
George Mason University
4400 University Drive
Science and Technology 2 MSN 4B5
Fairfax, VA 22030
12. Mr Steve Bridges
Technical Director
Joint Interoperability Test Command
P.O. Box 12798
Fort Huachuca, AZ 85670-2798
13. Dr Steve Hutchison
Director of Testing
Defense Information Systems Agency
Eagle Building
5111 Leesburg Pike
Falls Church, VA 22041-3205